



Mobile Intel[®] Pentium[®] II Processor Specification Update

Release Date: August 2001

Order Number: 243887-033

The Mobile Pentium[®] II processor may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are documented in this Specification Update.

Information in this document is provided in connection with Intel® products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel products are not intended for use in medical, life saving, or life sustaining applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The Mobile Pentium® II processor or the Intel® Pentium® II Processor Mobile Module may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

The Specification Update should be publicly available following the last shipment date for a period of time equal to the specific product's warranty period. Hardcopy Specification Updates will be available for one (1) year following End of Life (EOL). Web access will be available for three (3) years following EOL.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an ordering number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725 or by visiting Intel's website at <http://www.intel.com>

Copyright © Intel Corporation 1999, 2000, 2001.

Intel®, Pentium®, Celeron™, and Xeon™ are registered trademarks of Intel Corporation and its subsidiaries in the United States and other countries.

* Other brands and names may be claimed as the property of others.

CONTENTS

REVISION HISTORY.....	ii
PREFACE	iv
Specification Update for the Mobile Pentium® II Processor	1
GENERAL INFORMATION.....	1
Intel® Pentium® II Processor Mobile Module Markings	1
Intel® Pentium® II Processor Mobile Module with On-Die Cache Markings.....	4
Mobile Intel® Pentium® II Processor Mini-Cartridge Markings.....	5
Mobile Pentium® II Processor Mini-Cartridge Markings	6
Mobile Intel® Pentium® II Processor (Micro-PGA) Markings	7
Mobile Intel® Pentium® II Processor BGA Package Markings	7
IDENTIFICATION INFORMATION	8
SUMMARY OF CHANGES.....	12
Summary of Errata.....	13
Summary of Documentation Changes.....	16
Summary of Specification Clarifications	18
Summary of Specification Changes.....	19
ERRATA.....	20
DOCUMENTATION CHANGES.....	57
SPECIFICATION CLARIFICATIONS	76
SPECIFICATION CHANGES.....	81

REVISION HISTORY

Date of Revision	Version	Description
August 1998	-001	This document is the first Specification Update for the Mobile Pentium® II processor.
September 1998	-002	Updated Erratum 56. Added Errata 60 through 62.
October 1998	-003	Implemented new numbering nomenclature. Updated Errata B1 and B27. Added Specification Change B3. Added Documentation Change B21. Added Specification Clarifications B14 through B16. Added Errata B40 through B42.
November 1998	-004	Updated Erratum B23, Documentation Change B10 and B22. Added Erratum B43 and Specification Change B4.
December 1998	-005	Added Erratum B44. Updated Intel® Pentium® II Processor Mobile Module Identification Information table. Status update on Errata B13, B15, B17 through B20, B26, and B34.
January 1999	-006	Added Erratum B45, and Documentation Change B23 through B25.
January 25, 1999	-007	Special issue to coincide with launch of new processor. Added new product markings and updated the product ID tables. Updated Summary Table of Changes.
March 1999	-008	Added Specification Clarification B4 and Specification Change B4. Added S-Spec definition. Removed Specification Changes, Specification Clarifications, and Document Changes that have been incorporated into the appropriate documentation. Renumbered remaining items.
April 1999	-009	Added Erratum B46.
May 1999	-010	Added Erratum B47. Removed Document Changes that have been incorporated into the appropriate documentation and added Documentation Changes B5 through B9. Added Specification Clarifications B5 and B6. Added Specification Change B5. Updated BGA marking diagram.
June 1999	-011	Added µPGA1 marking diagram. Updated the Mobile Intel Pentium II Processor Identification Information table. Added Erratum B48. Added Documentation Changes B10 through B13. Added Specification Changes B6 and B7. Added mdp A0 stepping to all summary tables.
July 1999	-012	Updated Mobile Pentium II Processor Identification Information and Intel Pentium II Processor Mobile Module Identification Information tables. Added mqb A1 and mqp A1 stepping to all summary tables. Added Erratum B49. Added Documentation Change B14. Updated the Preface, Documentation Changes, Specification Clarifications and Specification Changes introduction paragraphs.
August 1999	-013	Updated Preface, Documentation Changes, Specification Clarifications, and Specification Changes paragraphs. Updated the

REVISION HISTORY

Date of Revision	Version	Description
		Identification Information table. Updated Intel Pentium II Processor Mobile Module at 266PE, 300PE, 333 and 366 MHz and 400 MHz Markings. Added Errata B50, B51. Added Documentation Change B15. Added Specification Clarification B7. Updated Codes Used in Summary Table.
October 1999	-014	Added erratum B52.
November 1999	-015	Added Errata B53 and B54; added Documentation Change B16, added Specification Clarification B8; added Brand ID column in the <i>Identification Information</i> table.
December 1999	-016	Added Erratum B55; added Documentation Change B17.
January 2000	-017	Added Erratum B56; added Documentation Change B18.
February 2000	-018	Revised Erratum B54; Added Erratum B57; Added Documentation Change B19. Updated Summary of Changes product letter codes.
March 2000	-019	Updated the Preface with new reference information. Revised Erratum B53; Added Erratum B58.
April 2000	-020	Updated the "Pentium II Processor Mobile Module Markings" and "Pentium II Processor Mobile Module With On-die Cache Markings" sections
June 2000	-021	Added Erratum B59.
July 2000	-022	Added Erratum B60, B61
August 2000	-023	Added Erratum B62
September 2000	-024	Added Erratum B63, B64; Revised Erratum B36, B53, B57; Added Documentation Changes B20, B21
October 2000	-025	Updated the reference to the published documents in the Preface; Added Erratum B65; Added Documentation Changes B22, B23
November 2000	-026	Added Erratum B66
December 2000	-027	Updated Specification Update product key to include the Intel® Pentium® 4 processor, Revised Erratum B2; Added Documentation Changes B24 thru B29.
January 2001	-028	Updated Summary of Changes; Revised Erratum B2; Added Documentation Changes B30, B31.
February 2001	-029	Revised Documentation Change B30 and added B32
March 2001	-030	Added Erratum B67 and B68
May 2001	-031	Added Erratum B69
July 2001	-032	Updated the "Pentium II Processor Mobile Module Identification Information"
August 2001	-033	Updated Summary of Changes; Added Erratum B70; Added Documentation Change B33

PREFACE

This document is an update to the specifications contained in the following documents:

- *Pentium® II Processor Developer's Manual* (Order Number 243341)
- *Mobile Pentium® II Processor at 233 MHz, 266 MHz, and 300 MHz* datasheet (Order Number 243669)
- *Intel® Pentium® II Processor Mobile Module: Mobile Module Connector 1 (MMC-1)* datasheet (Order Number 243667)
- *Intel® Pentium® II Processor Mobile Module: Mobile Module Connector 2 (MMC-2)* datasheet (Order Number 243668)
- *Mobile Pentium® II Processor in Mini-Cartridge Package at 400 MHz, 366 MHz, 333 MHz, 300PE MHz, and 266PE MHz* datasheet (Order Number 245108)
- *Mobile Pentium® II Processor in Micro-PGA and BGA Packages at 400 MHz, 366 MHz, 333 MHz, 300PE MHz, and 266PE MHz* datasheet (Order Number 245103)
- *Intel® Pentium® II Processor Mobile Module With On-Die Cache: Mobile Module Connector 1 (MMC-1)* datasheet (Order Number 245109)
- *Intel® Pentium® II Processor Mobile Module With On-Die Cache: Mobile Module Connector (MMC-2)* datasheet (Order Number 245110)
- *Intel Architecture Software Developer's Manual, Volumes 1, 2 and 3* (Order Numbers 243190, 243191, and 243192, respectively).
- *P6 Family of Processors Hardware Developer's Manual* (Order Number 244001)

It is intended for hardware system manufacturers and software developers of applications, operating systems, or tools. It contains S-Specs, Errata, Documentation Changes, Specification Clarifications and, Specification Changes.

Nomenclature

S-Spec Number is a five-digit code used to identify products. Products are differentiated by their unique characteristics, e.g., core speed, L2 cache size, package type, etc. as described in the processor identification information table. Care should be taken to read all notes associated with each S-Spec number.

Errata are design defects or errors. Errata may cause the Mobile Intel Celeron processor or the Intel Mobile Celeron Module's behavior to deviate from published specifications. Hardware and software designed to be used with any given processor must assume that all errata documented for that processor are present on all devices unless otherwise noted.

Documentation Changes include typos, errors, or omissions from the current published specifications. These changes will be incorporated in the next release of the specifications.

Specification Clarifications describe a specification in greater detail or further highlight a specification's impact to a complex design situation. These clarifications will be incorporated in the next release of the specifications.

Specification Changes are modifications to the current published specifications for the Intel® Mobile Celeron™ processor or the Intel® Celeron™ Processor Mobile Module. These changes will be incorporated in the next release of the specifications.

Specification Update for the Mobile Pentium® II Processor



GENERAL INFORMATION

Intel® Pentium® II Processor Mobile Module Markings

The Product Tracking Code (PTC) label shows the assembly level of the module and contains 13 characters. See the example below:

Example: ***PMD30005001AA***


AABCCCDDEEEFF

Definition:

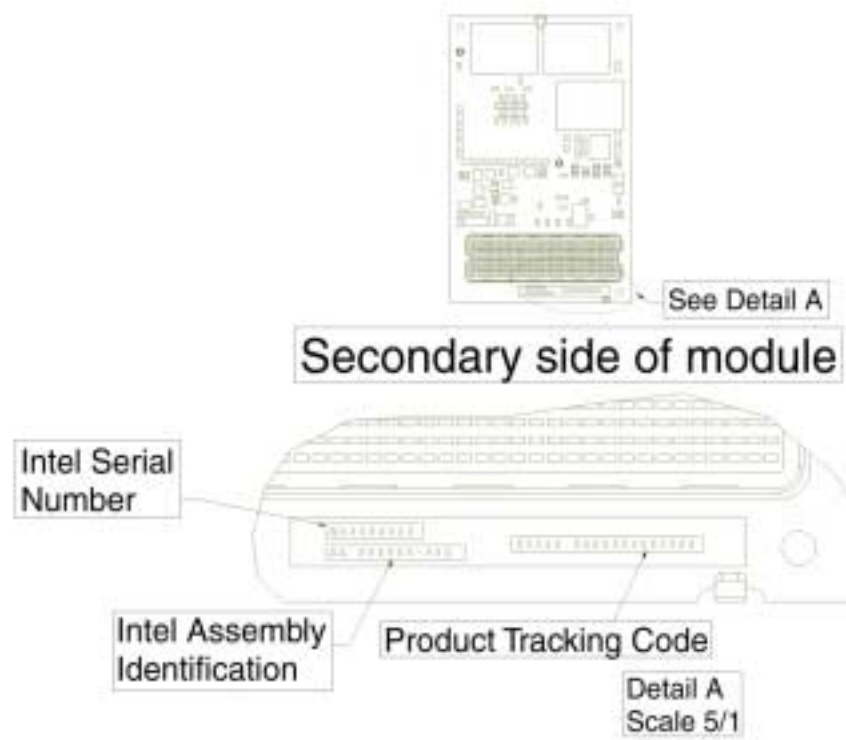
AA	Processor Module = PM
B	Pentium® II Processor Mobile Module (MMC-1) = D Pentium II Processor Mobile Module (MMC-2) = E
CCC	Speed Identity = 300, 266, 233
DD	Cache Size = 05 (512 KB)
EEE	Notifiable Design Revision (Start at 001)
FF	Notifiable Processor Revision (Start at AA)

Note: For other Intel Mobile Modules, the second field (B) is defined as:

Pentium® II Processor With On-die Cache Mobile Module (MMC-1) = F
Pentium® II Processor With On-die cache Mobile Module (MMC-2) = G
Celeron™ Processor Mobile Module (MMC-1) = H
Celeron™ Processor Mobile Module (MMC-2) = I



Pentium® II Processor Mobile Module (MMC-1)



Pentium® II Processor Mobile Module (MMC-2)

Intel® Pentium® II Processor Mobile Module with On-Die Cache Markings

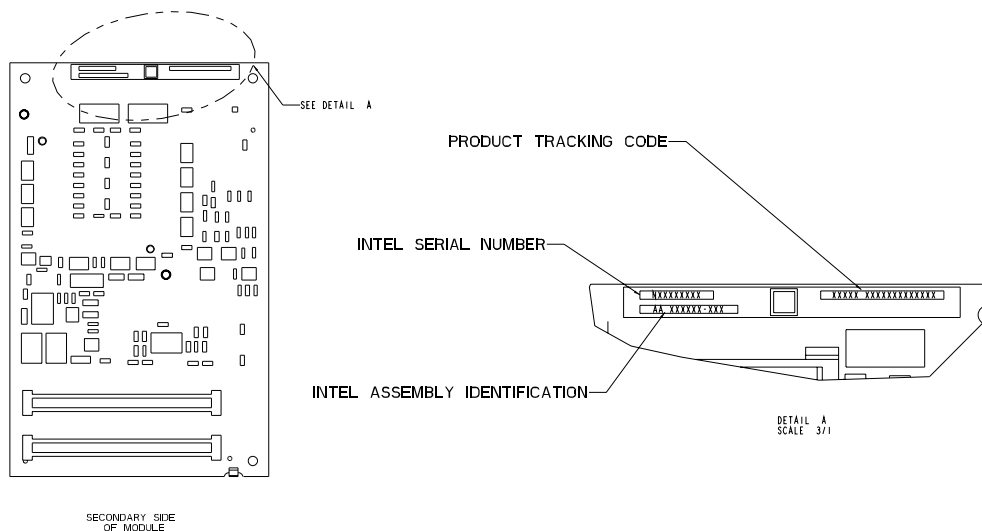
The Product Tracking Code (PTC) determines the Intel assembly level of the module. The PTC is on the secondary side of the module and contains 13 characters. See the example below:

Example: **PMG33302001AA**


Definition:	AABCCDDDEEFF	
	AA	Processor Module = PM
	B	Intel® Pentium® II Processor Mobile Module At 266PE, 300PE, 333, 366 and 400 MHz Markings Processor Mobile Module (MMC-1) = F Intel Pentium II Processor Mobile Module At 266PE, 300PE, 333, 366 and 400 MHz Markings Processor Mobile Module (MMC-2) = G
	CCC	Speed Identity = 400, 366, 333, 300, 266
	DD	Cache Size = 02 (256 KB)
	EEE	Notifiable Design Revision (Start at 001)
	FF	Notifiable Processor Revision (Start at AA)

Note: For other Intel Mobile Modules, the second field (B) is defined as:


Pentium® II Processor Mobile Module (MMC-1) = D
Pentium® II Processor Mobile Module (MMC-2) = E
Celeron™ Processor Mobile Module (MMC-1) = H
Celeron™ Processor Mobile Module (MMC-2) = I



Intel® Pentium® II Processor Mobile Module with On-Die Cache (MMC-1) Markings




Intel® Pentium® II Processor Mobile Module with On-Die Cache (MMC-2) Markings
Mobile Intel® Pentium® II Processor Mini-Cartridge Markings




Mobile Pentium® II Processor Mini-Cartridge Markings

Old Mobile Pentium® II Processor in Mini-Cartridge Package Label with Supplier Matrix Mark




New Mobile Pentium® II Processor in Mini-Cartridge Package Label with Supplier Matrix Mark Removed



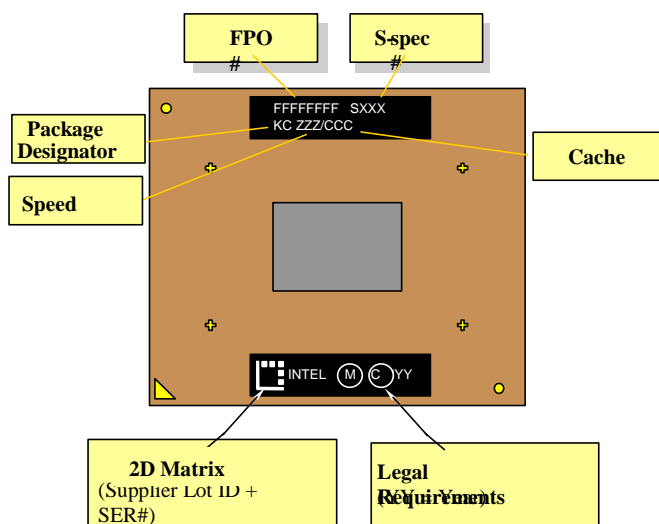
NOTES:

- FFFFFFFF = FPO # (Test Lot Traceability)
- NNNN = Serial Number
- XXXXX = Country of Origin
- SYYY = Production S-Spec
- ZZZ = Speed (MHz)
- LLL = Cache Size (k)

Mobile Intel® Pentium® II Processor (Micro-PGA) Markings



Mobile Intel® Pentium® II Processor BGA Package Markings



IDENTIFICATION INFORMATION

The Mobile Pentium® II processor or the Intel® Pentium® II Processor Mobile Module can be identified by the following values:

Family¹	233, 266, 300 Model 5²	233, 266, 300 Brand ID
0110	0101	00h (not supported)

Family¹	266PE, 300PE, 333, 366, 400 Model 6²	266PE, 300PE, 333, 366, 400 Brand ID
0110	0110	00h (not supported)

NOTES:

1. The Family corresponds to bits [11:8] of the EDX register after RESET, bits [11:8] of the EAX register after the CPUID instruction is executed with a 1 in the EAX register, and the generation field of the Device ID register accessible through Boundary Scan.
2. The Model corresponds to bits [7:4] of the EDX register after RESET, bits [7:4] of the EAX register after the CPUID instruction is executed with a 1 in the EAX register, and the model field of the Device ID register accessible through Boundary Scan.

The Mobile Pentium® II processor and the Intel® Pentium® II Processor Mobile Module's second level (L2) cache size can be determined by the following register contents:

512-Kbyte Unified L2 Cache¹	43h
256-Kbyte Unified L2 Cache	42h

NOTES:

- 1 For the Mobile Pentium® II processor and the Intel® Pentium® II Processor Mobile Module, the unified L2 cache size will be returned as one of the cache/TLB descriptors when the CPUID instruction is executed with a 2 in the EAX register.

Mobile Pentium® II Processor Identification Information

S-Spec	Core Steppings	CPUID	Speed (MHz) Core/Bus	L2 Size (Kbytes)	TagRAM/Stepping	Package	Notes
SL2KH	mdA0	0650h	233/66	512	T6P/A3	Mini-Cartridge	
SL2KJ	mdA0	0650h	266/66	512	T6P/A3	Mini-Cartridge	
SL2RS	mdB0	0652h	300/66	512	T6P-e/A0	Mini-Cartridge	1
SL2RR	mdB0	0652h	266/66	512	T6P-e/A0	Mini-Cartridge	1
SL2RQ	mdB0	0652h	233/66	512	T6P-e/A0	Mini-Cartridge	1
SL32M	mdxA0	066Ah	266/66	256	N/A	Mini-Cartridge	1
SL32N	mdxA0	066Ah	300/66	256	N/A	Mini-Cartridge	1
SL32P	mdxA0	066Ah	333/66	256	N/A	Mini-Cartridge	1
SL36Z	mdxA0	066Ah	366/66	256	N/A	Mini-Cartridge	1
SL3JW	MdxA0	066Ah	400/66	256	N/A	Mini-Cartridge	4
SL32Q	mdbA0	066Ah	266/66	256	N/A	BGA	2
SL32R	mdbA0	066Ah	300/66	256	N/A	BGA	2
SL32S	mdbA0	066Ah	333/66	256	N/A	BGA	2
SL3AG	mdbA0	066Ah	366/66	256	N/A	BGA	2
SL3EM	mqbA1	066Ah	400/66	256	N/A	BGA	5
SL3DR	mdbA0	066Ah	266/66	256	N/A	BGA	3
SL3HH	mdpA0	066Ah	266/66	256	N/A	Micro-PGA	2
SL3HJ	mdpA0	066Ah	300/66	256	N/A	Micro-PGA	2
SL3HK	mdpA0	066Ah	333/66	256	N/A	Micro-PGA	2
SL3HL	mdpA0	066Ah	366/66	256	N/A	Micro-PGA	2
SL3BW	mqpA1	066Ah	400/66	256	N/A	Micro-PGA	5

NOTES:

1. VCC_CORE is specified for 1.6 V +/-120 mV for these Intel® Mobile Pentium® II processors (mini-cartridge).
2. VCC_CORE is specified for 1.6 V +/-135 mV for these Mobile Intel Pentium II processors.
3. VCC_CORE is specified for 1.5 V +/-135 mV for these Mobile Intel Pentium II processors.
4. VCC_CORE is specified for 1.55 V +/-120 mV for these Mobile Intel Pentium II processors.
5. VCC_CORE is specified for 1.50 V +/-115 mV for these Mobile Intel Pentium II processors (0.18 µm).

Intel® Pentium® II Processor Mobile Module Identification Information

PTC	Product Steppings	CPUID	Speed (MHz) Core/Bus	L2 Size (Kbytes)	TagRAM/Stepping	Package	Notes
PMD23305001AA	mmdA0	0650h	233/66	512	T6P/A3	MMC-1	
PMD26605001AA	mmdA0	0650h	266/66	512	T6P/A3	MMC-1	
PMD30005002AA	mmdB0	0652h	300/66	512	T6P-e/A0	MMC-1	
PMD26605002AB	mmdB0	0652h	266/66	512	T6P-e/A0	MMC-1	
PMD23305002AB	mmdB0	0652h	233/66	512	T6P-e/A0	MMC-1	
PME23305001AA	mmdB0	0652h	233/66	512	T6P-e/A0	MMC-2	
PME26605001AA	mmdB0	0652h	266/66	512	T6P-e/A0	MMC-2	
PME30005001AA	mmdB0	0652h	300/66	512	T6P-e/A0	MMC-2	
PME30005002AA	mmdB0	0652h	300/66	512	T6P-e/A0	MMC-2	
PMG26602001AA	dmmA0	066Ah	266/66	256	N/A	MMC-2	1, 2, 3
PMG30002001AA	dmmA0	066Ah	300/66	256	N/A	MMC-2	1, 2, 3
PMG33302001AA	dmmA0	066Ah	333/66	256	N/A	MMC-2	1, 2, 3
PMG36602001AA	dmmA0	066Ah	366/66	256	N/A	MMC-2	1, 2, 3
PMG40002001AA	dmmA0	066Ah	400/66	256	N/A	MMC-2	1, 2, 4
PMF26602001AA	dmmA0	066Ah	266/66	256	N/A	MMC-1	1, 2, 3
PMF30002001AA	dmmA0	066Ah	300/66	256	N/A	MMC-1	1, 2, 3
PMF33302001AA	dmmA0	066Ah	333/66	256	N/A	MMC-1	1, 2, 3
PMF36602001AA	dmmA0	066Ah	366/66	256	N/A	MMC-1	1, 2, 3
PMF40002001AA	dmmA0	066Ah	400/66	256	N/A	MMC-1	1, 2, 4
PMG26602002AA	dmmA0	066Ah	266/66	256	N/A	MMC-2	1, 2, 4
PMG30002002AA	dmmA0	066Ah	300/66	256	N/A	MMC-2	1, 2, 4
PMG33302002AA	dmmA0	066Ah	333/66	256	N/A	MMC-2	1, 2, 4
PMG36602002AA	dmmA0	066Ah	366/66	256	N/A	MMC-2	1, 2, 4
PMG40002002AA	dmmA0	066Ah	400/66	256	N/A	MMC-2	1, 2, 5
PMF26602002AA	dmmA0	066Ah	266/66	256	N/A	MMC-1	1, 2, 4
PMF30002002AA	dmmA0	066Ah	300/66	256	N/A	MMC-1	1, 2, 4
PMF33302002AA	dmmA0	066Ah	333/66	256	N/A	MMC-1	1, 2, 4
PMF36602002AA	dmmA0	066Ah	366/66	256	N/A	MMC-1	1, 2, 4
PMF36602003AA	dmmA0	066Ah	366/66	256	N/A	MMC-1	1, 2, 5

NOTES:



MOBILE PENTIUM® II PROCESSOR SPECIFICATION UPDATE

1. V_{CC_CORE} is 1.6 V for these module products.
2. Voltage regulator modification implemented on these module products.
3. These are obsolete production units.
4. These are production units that have temp sensor modifications to improve module stability and reliability.
5. These are obsolete production units (with alternate AA) that have been relabeled and repackaged.

SUMMARY OF CHANGES

The following table indicates the Errata, Documentation Changes, Specification Clarifications, or Specification Changes that apply to Pentium® II processors. Intel intends to fix some of the errata in a future stepping of the component, and to account for the other outstanding issues through documentation or specification changes as noted. This table uses the following notations:

CODES USED IN SUMMARY TABLE

X:	Erratum, Documentation Change, Specification Clarification or Specification Change applies to the given processor stepping.
(No mark) or (blank box):	This item is fixed in or does not apply to the given stepping.
Fix:	This erratum is intended to be fixed in a future stepping of the component.
Fixed:	This erratum has been previously fixed.
NoFix:	There are no plans to fix this erratum.
Doc:	Intel intends to update the appropriate documentation in a future revision.
PKG:	This column refers to errata on the Mobile Pentium® II processor or the Intel® Pentium II Processor Mobile Module substrate.
AP:	APIC related erratum.
Shaded:	This item is either new or modified from the previous version of the document.

Each Specification Update item is prefixed with a capital letter to distinguish the product. The key below details the letters that are used in Intel's microprocessor Specification Updates:

A = Intel® Pentium® II processor

B = Mobile Intel® Pentium® II processor

C = Intel® Celeron™ processor

D = Intel® Pentium® II Xeon™ processor

E = Intel® Pentium® III processor

G = Intel® Pentium® III Xeon™ processor

H = Mobile Intel® Celeron™ processor at 466 MHz, 433 MHz, 400 MHz, 366 MHz, 333 MHz, 300 MHz, and 266 MHz

K = Mobile Intel® Pentium® III processor

M = Mobile Intel® Celeron™ processor at 500 MHz, 450 MHz, and 400A MHz

N = Intel® Pentium® 4 processor

P = Intel® Xeon™ processor

The Specification Updates for the Pentium® processor, Pentium® Pro processor, and other Intel products do not use this convention.

Summary of Errata

NO.	m dA0	mmd A0	m dB0	mmd B0	mdx A0	mdb A0	dmm A0	mqb A1	mdp A0	mqp A1	Plans	ERRATA
B1	X	X	X	X	X	X	X	X	X	X	NoFix	FP Data Operand Pointer may be incorrectly calculated after FP access which wraps 64-Kbyte boundary in 16-bit code
B2	X	X	X	X	X	X	X	X	X	X	NoFix	Differences exist in debug exception reporting
B3	X	X	X	X	X	X	X	X	X	X	NoFix	Code fetch matching disabled debug register may cause debug exception
B4	X	X	X	X	X	X	X	X	X	X	NoFix	Double ECC error on read may result in BINIT#
B5	X	X	X	X	X	X	X	X	X	X	NoFix	FP inexact-result exception flag may not be set
B6	X	X	X	X	X	X	X	X	X	X	NoFix	BTM for SMI will contain incorrect FROM EIP
B7	X	X	X	X	X	X	X	X	X	X	NoFix	I/O restart in SMM may fail after simultaneous MCE
B8	X	X	X	X	X	X	X	X	X	X	NoFix	Branch traps do not function if BTMs are also enabled
B9	X	X	X	X	X	X	X	X	X	X	NoFix	Machine check exception handler may not always execute successfully
B10	X	X	X	X	X	X	X	X	X	X	NoFix	MCE due to L2 parity error gives L1 MCACOD.LL
B11	X	X	X	X	X	X	X	X	X	X	NoFix	LBERR may be corrupted after some events
B12	X	X	X	X	X	X	X	X	X	X	NoFix	BTMs may be corrupted during simultaneous L1 cache line replacement
B13	X	X			X	X	X	X	X	X	Fixed	Potential early deassertion of LOCK# during split-lock cycles
B14	X	X	X	X	X	X	X	X	X	X	NoFix	A20M# may be inverted after returning from SMM and Reset
B15	X	X			X	X	X	X	X	X	Fix	Reporting of floating-point exception may be delayed
B16	X	X	X	X	X	X	X	X	X	X	NoFix	Near CALL to ESP creates unexpected EIP address
B17	X	X			X	X	X	X	X	X	Fix	Built-in self test always gives nonzero result
B18	X	X			X	X	X	X	X	X	Fix	Cache state corruption in the presence of page A/D-bit setting and snoop traffic
B19	X	X			X	X	X	X	X	X	Fix	Snoop cycle generates spurious machine check exception
B20	X	X			X	X	X	X	X	X	Fix	MOVD/MOVQ instruction writes to memory prematurely
B21	X	X	X	X	X	X	X	X	X	X	NoFix	Memory type undefined for nonmemory operations

Summary of Errata

NO.	m dA0	mmd A0	m dB0	mmd B0	mdx A0	mdb A0	dmm A0	mqb A1	mdp A0	mqp A1	Plans	ERRATA
B22	X	X	X	X	X	X	X	X	X	X	NoFix	FP Data Operand Pointer may not be zero after power on or Reset
B23	X	X	X	X	X	X	X	X	X	X	NoFix	MOVD following zeroing instruction can cause incorrect result
B24	X	X	X	X	X	X	X	X	X	X	NoFix	Premature execution of a load operation prior to exception handler invocation
B25	X	X	X	X	X	X	X	X	X	X	NoFix	Read portion of RMW instruction may execute twice
B26	X		X								Fixed	Voltage sequencing during power up
B27	X	X	X	X	X	X	X	X	X	X	Fix	Intervening writeback may occur during locked transaction
B28	X	X	X	X	X	X	X	X	X	X	NoFix	MC2_STATUS MSR has model-specific error code and machine check architecture error code reversed
B29	X	X	X	X	X	X	X	X	X	X	NoFix	Mixed cacheability of lock variables is problematic in MP systems
B30	X	X	X	X	X	X	X	X	X	X	Fix	Thermal sensor may assert SMBALERT# incorrectly
B31	X	X	X	X	X	X	X	X	X	X	NoFix	MOV with debug register causes debug exception
B32	X	X	X	X	X	X	X	X	X	X	NoFix	Upper four PAT entries not usable with Mode B or Mode C paging
B33	X	X	X	X							Fix	UC write may be reordered around a cacheable write
B34	X	X			X	X	X	X	X	X	Fix	Incorrect memory type may be used when MTRRs are disabled
B35	X	X	X	X							Fix	Misprediction in program flow may cause unexpected instruction execution
B36	X	X	X	X	X	X	X	X	X	X	NoFix	Data breakpoint exception in a displacement relative near call may corrupt EIP
B37	X	X	X	X	X	X	X	X	X	X	NoFix	System bus ECC not functional with 2:1 ratio
B38	X	X	X	X							Fix	Fault on REP CMPS/SCAS operation may cause incorrect EIP
B39	X	X	X	X	X	X	X	X	X	X	NoFix	RDMSR and WRMSR to invalid MSR may not cause GP fault
B40	X	X	X	X	X	X	X	X	X	X	NoFix	SYSENTER/SYSEXIT instructions can implicitly load "null segment selector" to SS and CS registers
B41	X	X	X	X	X	X	X	X	X	X	NoFix	PRELOAD followed by EXTEST does not load boundary scan data
B42	X	X	X	X	X	X	X	X	X	X	NoFix	Far jump to new TSS with D-bit cleared may cause system hang
B43	X	X	X	X	X	X	X	X	X	X	NoFix	Incorrect chunk ordering may prevent execution of the Machine

Summary of Errata

NO.	m dA0	mmd A0	m dB0	mmd B0	mdx A0	mdb A0	dmm A0	mqb A1	mdp A0	mqp A1	Plans	ERRATA
												Check Exception handler after BINIT#
B44	X	X	X	X	X	X	X	X	X	X	Fix	Resume flag may not be cleared after debug exception
B45		X		X							Fixed	Thermal sensor leakage current may exceed specification
B46							X				Fix	Processor may return invalid parameters on execution of the CPUID instruction
B47					X	X	X	X	X	X	NoFix	Internal cache protocol violation may cause system hang
B48	X	X	X	X	X	X	X	X	X	X	NoFix	GP# fault on WSMSR to ROB_CR_BKUPTMPDR6
B49	X	X	X	X	X	X	X	X	X	X	NoFix	Machine Check Exception may occur due to improper line eviction in the IFU
B50								X		X	NoFix	System may hang when the Mobile Pentium® II processor is in the Deep Sleep state
B51						X			X		NoFix	VTOL pin may be driven to ground on the Mobile Pentium II processor
B52	X	X	X	X	X	X	X	X	X	X	NoFix	Lower Bits of SMRAM SMBASE Register Cannot Be Written With an ITP
B53	X	X	X	X	X	X	X	X	X	X	NoFix	Task switch may cause wrong PTE and PDE access bit to be set
B54	X	X	X	X	X	X	X	X	X	X	NoFix	Unsynchronized cross-modifying code operations can cause unexpected instruction execution results
B55	X	X	X	X	X	X	X	X	X	X	NoFix	Deadlock may occur due to illegal-instruction/page-miss Combination
B56	X	X	X	X	X	X	X	X	X	X	NoFix	FLUSH# assertion following STPCLK# may prevent CPU clocks from stopping
B57	X	X	X	X	X	X	X	X	X	X	NoFix	Floating-point exception condition can be deferred
B58		X		X			X				NoFix	Race conditions may exist on thermal sensor SMBus collision detection/arbitration circuitry
B59	X	X	X	X	X	X	X	X	X	X	NoFix	Intermittent power-on failure due to uninitialized processor internal nodes
B60	X	X	X	X	X	X	X	X	X	X	NoFix	Selector for the LTR/LLDT register may get corrupted
B61	X	X	X	X	X	X	X	X	X	X	NoFix	INIT does not clear global entries in the TLB
B62	X	X	X	X	X	X	X	X	X	X	NoFix	VM bit will be cleared on a double fault handler
B63	X	X	X	X	X	X	X	X	X	X	NoFix	Memory aliasing with inconsistent A and D bits may cause processor

Summary of Errata

NO.	m dA0	mmd A0	m dB0	mmd B0	mdx A0	mdb A0	dmm A0	mqb A1	mdp A0	mqp A1	Plans	ERRATA
												deadlock
B64	X	x	X	X	X	X	X	X	X	X	NoFix	Use of memory aliasing with inconsistent memory type may cause system hang
B65	X	X	X	X	X	X	X	X	X	X	NoFix	Processor may report invalid TSS fault instead of Double fault during mode C paging
B66	X	X	X	X	X	X	X	X	X	X	NoFix	Machine check exception may occur when interleaving code between different memory types
B67	X	X	X	X	X	X	X	X	X	X	NoFix	Wrong ESP Register Values During a Fault in VM86 Mode
B68	X	X	X	X	X	X	X	X	X	X	NoFix	APIC ICR Write May Cause Interrupt Not to be Sent When ICR Delivery Bit Pending
B1AP	X	X	X	X	X	X	X	X	X	X	NoFix	APIC access to cacheable memory causes SHUTDOWN
B69	X	X	X	X	X	X	X	X	X	X	NoFix	Processor Incorrectly Samples NMI Interrupt after RESET# Deassertion When Processor APIC is Hardware-Disabled
B70	X	X	X	X	X	X	X	X	X	X	NoFix	The Instruction Fetch Unit (IFU) May Fetch Instructions Based Upon Stale CR3 Data After a Write to CR3 Register

Summary of Documentation Changes

NO.	m dA0	mmd A0	m dB0	mmd B0	mdx A0	mdb A0	dmm A0	mqb A1	mdp A0	mqp A1	Plans	Documentation Changes
B1	X		X								Doc	Mislabeled vertical dimension parameters
B2	X		X								Doc	Correction to dimension definitions
B3	X		X									Correction to vertical dimensions
B4	X	X	X	X	X	X	X	X	X	X	Doc	SMBus data setup time
B5											Doc	Intel® Pentium® II Processor Mobile Module Connector 1 (MMC-1) northbridge correction
B6							X				Doc	VR_ON and VRPWRGD specification correction for MMC-1 and MMC-2

Summary of Documentation Changes

NO.	m dA0	mmd A0	m dB0	mmd B0	mdx A0	mdb A0	dmm A0	mqb A1	mdp A0	mqp A1	Plans	Documentation Changes
B7							X				Doc	STR and POS measurement correction for MMC-1 and MMC-2
B8							X				Doc	L2 cache correction for MMC-1 and MMC-2
B9					X	X		X	X	X	Doc	Die height (A ₂) in surface-mount BGA1 package specification
B10	X	X	X	X	X	X	X	X	X	X	Doc	MTRR initialization clarification
B11	X	X	X	X	X	X	X	X	X	X	Doc	STPCLK# pin definition
B12	X	X	X	X	X	X	X	X	X	X	Doc	RESET# pin definition
B13						X		X			Doc	BGA1 package mechanical specifications
B14					X	X		X	X	X	Doc	Generic clock waveform figure
B15	X	X	X	X	X	X	X	X	X	X	Doc	Invalidating caches and TLBs
B16	X	X	X	X	X	X	X	X	X	X	Doc	Handling of self-modifying and cross-modifying code
B17	X	X	X	X	X	X	X	X	X	X	Doc	Machine Check Architecture initialization of MCI_STATUS registers
B18	X	X	X	X	X	X	X	X	X	X	Doc	LOCK# signal prefix operands
B19	X	X	X	X	X	X	X	X	X	X	Doc	SMRAM state save map contains documentation errors
B20	X	X	X	X	X	X	X	X	X	X	Doc	System Management Interrupt (SMI) during startup IPI clarification
B21	X	X	X	X	X	X	X	X	X	X	Doc	Memory aliasing with different memory types
B22	X	X	X	X	X	X	X	X	X	X	Doc	Runbist will not function when stpclk# driven low
B23	X	X	X	X	X	X	X	X	X	X	Doc	Memory aliasing with inconsistent A and D bits may cause processor deadlock
B24	X	X	X	X	X	X	X	X	X	X	Doc	An interrupt could occur while TSS is marked busy
B25	X	X	X	X	X	X	X	X	X	X	Doc	NMI unmasked early when processor is running in V86 mode

Summary of Documentation Changes

NO.	m dA0	mmd A0	m dB0	mmd B0	mdx A0	mdb A0	dmm A0	mqb A1	mdp A0	mqp A1	Plans	Documentation Changes
B26	X	X	X	X	X	X	X	X	X	X	Doc	P6 family processors read two bytes for POP SEG instruction
B27	X	X	X	X	X	X	X	X	X	X	Doc	APIC register offsets are aligned on 128-bit boundaries
B28	X	X	X	X	X	X	X	X	X	X	Doc	Single stepping of instructions breaks out of HLT state
B29	X	X	X	X	X	X	X	X	X	X	Doc	Additional signal resumes execution while in a HALT state
B30	X	X	X	X	X	X	X	X	X	X	Doc	PDPTLRloads are always uncacheable
B31	X	X	X	X	X	X	X	X	X	X	Doc	SMI during HALT causes PMC miscounts of retired instructions
B32	X	X	X	X	X	X	X	X	X	X	Doc	INIT Message is Not Sent Out Through APIC
B33	X	X	X	X	X	X	X	X	X	X	Doc	Mobile Pentium II Processor CPUID Section Update

Summary of Specification Clarifications

NO.	m dA0	mmd A0	m dB0	mmd B0	mdx A0	mdb A0	dmm A0	mqb A1	mdp A0	mqp A1	Plans	Specification Clarifications
B1	X	X	X	X	X	X	X	X	X	X	Doc	PWRGOOD inactive pulse width
B2	X	X	X	X	X	X	X	X	X	X	Doc	Thermal sensor setpoint
B3	X	X	X	X	X	X	X	X	X	X	Doc	Thermal sensor configuration register – RUN/STOP bit
B4		X		X								Thermal attach point gageable thread clarification
B5						X		X	X	X	Doc	Shipping container maximum temperature rating for BGA1
B6							X				Doc	ESD for MMC-1 and MMC-2
B7								X		X	Doc	Thermal diode footnote
B8		X		X			X				Doc	Bulk capacitance requirements for the system electronics

Summary of Specification Changes

NO.	m dA0	mmd A0	m dB0	mmd B0	mdx A0	mdb A0	dmm A0	mqb A1	mdp A0	mqp A1	Plans	Specification Changes
B1	X		X								Doc	Maximum die and cover pressure specifications
B2	X	X	X	X	X	X	X	X	X	X	Doc	New footnote for PWRGOOD inactive pulse width
B3	X	X	X	X	X	X	X	X	X	X	Doc	APIC reference correction
B4	X		X									Low Power GTL+ Output Leakage Current change
B5					X	X		X	X	X	Doc	I _{CC,SG} , I _{CC,QS} , and I _{CC,DSL} maximum specifications for BGA1 and mini-cartridge
B6						X		X	X	X		VTOL signal definition and its voltage tolerance
B7					X	X		X	X	X	Doc	Temperature note for thermal power specifications

ERRATA

B1. FP Data Operand Pointer May Be Incorrectly Calculated After FP Access Which Wraps 64-Kbyte Boundary in 16-Bit Code

Problem: The FP Data Operand Pointer is the effective address of the operand associated with the last noncontrol floating-point instruction executed by the machine. If an 80-bit floating-point access (load or store) occurs in a 16-bit mode other than protected mode (in which case the access will produce a segment limit violation), the memory access wraps a 64-Kbyte boundary, and the floating-point environment is subsequently saved, the value contained in the FP Data Operand Pointer may be incorrect.

Implication: A 32-bit operating system running 16-bit floating-point code may encounter this erratum, under the following conditions:

- The operating system is using a segment greater than 64 Kbytes in size.
- An application is running in a 16-bit mode other than protected mode.
- An 80-bit floating-point load or store which wraps the 64-Kbyte boundary is executed.
- The operating system performs a floating-point environment store (FSAVE/FNSAVE/FSTENV/FNSTENV) after the above memory access.
- The operating system uses the value contained in the FP Data Operand Pointer.

Wrapping an 80-bit floating-point load around a segment boundary in this way is not a normal programming practice. Intel has not currently identified any software which exhibits this behavior.

Workaround: If the FP Data Operand Pointer is used in an OS which may run 16-bit floating-point code, care must be taken to ensure that no 80-bit floating-point accesses are wrapped around a 64-Kbyte boundary.

Status: For the steppings affected see the Summary Table of Changes at the beginning of this section.

B2. Differences Exist in Debug Exception Reporting

Problem: There exist some differences in the reporting of code and data breakpoint matches between that specified by previous Intel processor specifications and the behavior of the Intel® Mobile Pentium® II processor, as described below:

Case 1: The first case is for a breakpoint set on a MOVSS or POPSS instruction, when the instruction following it causes a debug register protection fault (DR7.gd is already set, enabling the fault). The processor reports delayed data breakpoint matches from the MOVSS or POPSS instructions by setting the matching DR6.bi bits, along with the debug register protection fault (DR6.bd). If additional breakpoint faults are matched during the call of the debug fault handler, the processor sets the breakpoint match bits (DR6.bi) to reflect the breakpoints matched by both the MOVSS or POPSS breakpoint and the debug fault handler call. The Mobile Intel Pentium II processor only sets DR6.bd in either situation, and does not set any of the DR6.bi bits.

Case 2: In the second breakpoint reporting failure case, if a MOVSS or POPSS instruction with a data breakpoint is followed by a store to memory which:

a) crosses a 4-Kbyte page boundary,

OR

b) causes the page table Access or Dirty (A/D) bits to be modified,

the breakpoint information for the MOVSS or POPSS will be lost. Previous processors retain this information under these boundary conditions.

Case 3: If they occur after a MOVSS or POPSS instruction, the INTn, INTO, and INT3 instructions zero the DR6.bi bits (bits B0 through B3), clearing pending breakpoint information, unlike previous processors.

Case 4: If a data breakpoint and an SMI (System Management Interrupt) occur simultaneously, the SMI will be serviced via a call to the SMM handler, and the pending breakpoint will be lost.

Case 5: When an instruction that accesses a debug register is executed, and a breakpoint is encountered on the instruction, the breakpoint is reported twice.

Case 6: Unlike previous versions of Intel Architecture processors, Mobile Intel Pentium II processors will not set the Bi bits for a matching disabled breakpoint unless at least one other breakpoint is enabled.

Implication: When debugging or when developing debuggers for a Mobile Intel Pentium II processor-based system, this behavior should be noted. Normal usage of the MOVSS or POPSS instructions (i.e., following them with a MOV ESP) will not exhibit the behavior of cases 1-3. Debugging in conjunction with SMM will be limited by case 4.

Workaround: Following MOVSS and POPSS instructions with a MOV ESP instruction when using breakpoints will avoid the first three cases of this erratum. No workaround has been identified for cases 4, 5, or 6.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B3. Code Fetch Matching Disabled Debug Register May Cause Debug Exception

Problem: The bits L0-3 and G0-3 enable breakpoints local to a task and global to all tasks, respectively. If one of these bits is set, a breakpoint is enabled, corresponding to the addresses in the debug registers DR0-DR3. If at least one of these breakpoints is enabled, any of these registers are *disabled* (e.g., *Ln* and *Gn* are 0), and *RWn* for the disabled register is 00 (indicating a breakpoint on instruction execution), normally an instruction fetch will not cause an instruction-breakpoint fault based on a match with the address in the disabled register(s). However, if the address in a disabled register matches the address of a code fetch which also results in a page fault, an instruction-breakpoint fault will occur.

Implication: While debugging software, extraneous instruction-breakpoint faults may be encountered if breakpoint registers are not cleared when they are disabled. Debug software which does not implement a code breakpoint handler will fail, if this occurs. If a handler is present, the fault will be serviced. Mixing data and code may exacerbate this problem by allowing disabled data breakpoint registers to break on an instruction fetch.

Workaround: The debug handler should clear breakpoint registers before they become disabled.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B4. Double ECC Error on Read May Result in BINIT#

Problem: For this erratum to occur, the following conditions must be met:

- Machine Check Exceptions (MCEs) must be enabled.
- A dataless transaction (such as a write invalidate) must be occurring simultaneously with a transaction which returns data (a normal read).
- The read data must contain a double-bit uncorrectable ECC error.

If these conditions are met, the Mobile Pentium II processor will not be able to determine which transaction was erroneous, and instead of generating an MCE, it will generate a BINIT#.

Implication: The bus will be reinitialized in this case. However, since a double-bit uncorrectable ECC error occurred on the read, the MCE handler (which is normally reached on a double-bit uncorrectable ECC error for a read) would most likely cause the same BINIT# event.

Workaround: Though the ability to drive BINIT# can be disabled in the Mobile Pentium II processor, which would prevent the effects of this erratum, overall system behavior would not improve, since the error which would normally cause a BINIT# would instead cause the machine to shut down. No other workaround has been identified.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B5. *FP Inexact-Result Exception Flag May Not Be Set*

Problem: When the result of a floating-point operation is not exactly representable in the destination format (1/3 in binary form, for example), an inexact-result (precision) exception occurs. When this occurs, the PE bit (bit 5 of the FPU status word) is normally set by the processor. Under certain rare conditions, this bit may not be set when this rounding occurs. However, other actions taken by the processor (invoking the software exception handler if the exception is unmasked) are not affected. This erratum can only occur if the floating-point operation which causes the precision exception is immediately followed by one of the following instructions:

- FST m32real
- FST m64real
- FSTP m32real
- FSTP m64real
- FSTP m80real
- FIST m16int
- FIST m32int
- FISTP m16int
- FISTP m32int
- FISTP m64int

Note that even if this combination of instructions is encountered, there is also a dependency on the internal pipelining and execution state of both instructions in the processor.

Implication: Inexact-result exceptions are commonly masked or ignored by applications, as it happens frequently, and produces a rounded result acceptable to most applications. The PE bit of the FPU status word may not always be set upon receiving an inexact-result exception. Thus, if these exceptions are unmasked, a floating-point error exception handler may not recognize that a precision exception occurred. Note that this is a “sticky” bit, e.g., once set by an inexact-result condition, it remains set until cleared by software.

Workaround: This condition can be avoided by inserting two NOP instructions between the two floating-point instructions.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B6. *BTM for SMI Will Contain Incorrect FROM EIP*

Problem: A system management interrupt (SMI) will produce a Branch Trace Message (BTM), if BTMs are enabled. However, the FROM EIP field of the BTM (used to determine the address of the instruction which was being executed when the SMI was serviced) will not have been updated for the SMI, so the field will report the same FROM EIP as the previous BTM.

Implication: A BTM which is issued for an SMI will not contain the correct FROM EIP, limiting the usefulness of BTMs for debugging software in conjunction with System Management Mode (SMM).

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B7. I/O Restart in SMM May Fail After Simultaneous MCE

Problem: If an I/O instruction (IN, INS, REP INS, OUT, OUTS, or REP OUTS) is being executed, and if the data for this instruction becomes corrupted, the Mobile Pentium II processor will signal a machine check exception (MCE). If the instruction is directed at a device which is powered down, the processor may also receive an assertion of SMI#. Since MCEs have higher priority, the processor will call the MCE handler, and the SMI# assertion will remain pending. However, upon attempting to execute the first instruction of the MCE handler, the SMI# will be recognized and the processor will attempt to execute the SMM handler. If the SMM handler is completed successfully, it will attempt to restart the I/O instruction, but will not have the correct machine state, due to the call to the MCE handler.

Implication: A simultaneous MCE and SMI# assertion may occur for one of the I/O instructions above. The SMM handler may attempt to restart such an I/O instruction, but will have corrupted state due to the MCE handler call, leading to failure of the restart and shutdown of the processor.

Workaround: If a system implementation must support both SMM and MCEs, the first thing the SMM handler code (when an I/O restart is to be performed) should do is check for a pending MCE. If there is an MCE pending, the SMM handler should immediately exit via an RSM instruction and allow the machine check exception handler to execute. If there is not, the SMM handler may proceed with its normal operation.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B8. Branch Traps Do Not Function If BTMs Are Also Enabled

Problem: If branch traps or branch trace messages (BTMs) are enabled alone, both function as expected. However, if both are enabled, only the BTMs will function, and the branch traps will be ignored.

Implication: The branch traps and branch trace message debugging features cannot be used together.

Workaround: If branch trap functionality is desired, BTMs must be disabled.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B9. Machine Check Exception Handler May Not Always Execute Successfully

Problem: An asynchronous machine check exception (MCE), such as a BINIT# event, which occurs during an access that splits a 4-Kbyte page boundary may leave some internal registers in an indeterminate state. Thus, MCE handler code may not always run successfully if an asynchronous MCE has occurred previously.

Implication: An MCE may not always result in the successful execution of the MCE handler. However, asynchronous MCEs usually occur upon detection of a catastrophic system condition that would also hang the processor. Leaving MCEs disabled will result in the condition which caused the asynchronous MCE instead causing the processor to enter shutdown. Therefore, leaving MCEs disabled may not improve overall system behavior.

Workaround: No workaround which would guarantee successful MCE handler execution under this condition has been identified.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B10. MCE Due to L2 Parity Error Gives L1 MCACOD.LL

Problem: If a Cache Reply Parity (CRP) error, Cache Address Parity (CAP) error, or Cache Synchronous Error (CSER) occurs on an access to the Mobile Pentium II processor's L2 cache, the resulting Machine Check Architectural Error Code (MCACOD) will be logged with '01' in the LL field. This value indicates an L1 cache error; the value should be '10', indicating an L2 cache error. Note that L2 ECC errors have the correct value of '10' logged.

Implication: An L2 cache access error, other than an ECC error, will be improperly logged as an L1 cache error in MCACOD.LL.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B11. LBER May Be Corrupted After Some Events

Problem: The last branch record (LBR) and the last branch before exception record (LBER) can be used to determine the source and destination information for previous branches or exceptions. The LBR contains the source and destination addresses for the last branch or exception, and the LBER contains similar information for the last branch taken before the last exception. This information is typically used to determine the location of a branch which leads to execution of code which causes an exception. However, after a catastrophic bus condition which results in an assertion of BINIT# and the reinitialization of the buses, the value in the LBER may be corrupted. Also, after either a CALL which results in a fault or a software interrupt, the LBER and LBR will be updated to the same value, when the LBER should not have been updated.

Implication: The LBER and LBR registers are used only for debugging purposes. When this erratum occurs, the LBER will not contain reliable address information. The value of LBER should be used with caution when debugging branching code; if the values in the LBR and LBER are the same, then the LBER value is incorrect. Also, the value in the LBER should not be relied upon after a BINIT# event.

Workaround: None identified at this time.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B12. BTMs May Be Corrupted During Simultaneous L1 Cache Line Replacement

Problem: When Branch Trace Messages (BTMs) are enabled and such a message is generated, the BTM may be corrupted when issued to the bus by the L1 cache if a new line of data is brought into the L1 data cache simultaneously. Though the new line being stored in the L1 cache is stored correctly, and no corruption occurs in the data, the information in the BTM may be incorrect due to the internal collision of the data line and the BTM.

Implication: Although BTMs may not be entirely reliable due to this erratum, the conditions necessary for this boundary condition to occur have only been exhibited during focused simulation testing. Intel has currently not observed this erratum in a system level validation environment.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B13. Potential Early Deassertion of LOCK# During Split-Lock Cycles

Problem: During a split-lock cycle there are four bus transactions: 1st ADS# (a partial read), 2nd ADS# (a partial read), 3rd ADS# (a partial write), and the 4th ADS# (a partial write). Due to this erratum, LOCK# may deassert one clock after the 4th ADS# of the split-lock cycle instead of after the 4th RS# assertion corresponding to the 4th ADS# has been sampled. The following sequence of events are required for this erratum to occur:

1. A lock cycle occurs (split or nonsplit).
2. Five more bus transactions (assertion of ADS#) occur.
3. A split-lock cycle occurs and BNR# toggles after the 3rd ADS# (partial write) of the split-lock cycle. This in turn delays the assertion of the 4th ADS# of the split-lock cycle. BNR# toggling at this time could most likely happen when the bus is set for an IOQ depth of 2.

When all of these events occur, LOCK# will be deasserted in the next clock after the 4th ADS# of the split-lock cycle.

Implication: This may affect chipset logic which monitors the behavior of LOCK# deassertion.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B14. A20M# May Be Inverted After Returning From SMM and Reset

Problem: This erratum is seen when software causes the following events to occur:

1. The assertion of A20M# in real address mode.
2. After entering the 1-Mbyte address wrap-around mode caused by the assertion of A20M#, there is an assertion of SMI# intended to cause a Reset or remove power to the processor. Once in the SMM handler, software saves the SMM state save map to an area of nonvolatile memory from which it can be restored at some point in the future. Then software asserts RESET# or removes power to the processor.
3. After exiting Reset or completion of power-on, software asserts SMI# again. Once in the SMM handler, it then retrieves the old SMM state save map which was saved in event 2 above and copies it into the current SMM state save map. Software then asserts A20M# and executes the RSM instruction. After exiting the SMM handler, the polarity of A20M# is inverted.

Implication: If this erratum occurs, A20M# will behave with a polarity opposite from what is expected (e.g., the 1-Mbyte address wrap-around mode is enabled when A20M# is deasserted, and does not occur when A20M# is asserted).

Workaround: Software should save the A20M# signal state in nonvolatile memory before an assertion of RESET# or a power down condition. After coming out of Reset or at power on, SMI# should be asserted again. During the restoration of the old SMM state save map described in event 3 above, the entire map should be restored, except for bit 5 of the byte at offset 7F18h. This bit should retain the value assigned to it when the SMM state save map was created in event 3. The SMM handler should then restore the original value of the A20M# signal.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B15. Reporting of Floating-Point Exception May Be Delayed

Problem: The Mobile Pentium II processor normally reports a floating-point exception for an instruction when the next floating-point or MMX™ technology instruction is executed. The assertion of FERR# and/or the INT 16 interrupt corresponding to the exception may be delayed until the floating-point or MMX technology instruction *after* the one which is expected to trigger the exception, if the following conditions are met:

1. A floating-point instruction causes an exception.
2. Before another floating-point or MMX technology instruction, any one of the following occurs:
 - a. A subsequent data access occurs to a page which has not been marked as accessed, or
 - b. Data is referenced which crosses a page boundary, or
 - c. A possible page-fault condition is detected which, when resolved, completes without faulting.
3. The instruction causing event 2 above is followed by a MOVQ or MOVD store instruction.

Implication: This erratum only affects software which operates with floating-point exceptions unmasked. Software which requires floating-point exceptions to be visible on the next floating-point or MMX technology instruction, and which uses floating-point calculations on data which is then used for MMX technology instructions, may see a delay in the reporting of a floating-point instruction exception in some cases. Note that mixing floating-point and MMX technology instructions in this way is not recommended.

Workaround: Inserting a WAIT or FWAIT instruction (or reading the floating-point status register) between the floating-point instruction and the MOVQ or MOVD instruction will give the expected results. This is already the recommended practice for software.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B16. Near CALL to ESP Creates Unexpected EIP Address

Problem: As documented, the CALL instruction saves procedure linking information in the procedure stack and jumps to the called procedure specified with the destination (target) operand. The target operand specifies the address of the first instruction in the called procedure. This operand can be an immediate value, a general purpose register, or a memory location. When accessing an absolute address indirectly using the stack pointer (ESP) as a base register, the base value used is the value in the ESP register before the instruction executes. However, when accessing an absolute address directly using ESP as the base register, the base value used is the value of ESP *after* the return value is pushed on the stack, not the value in the ESP register *before* the instruction executed.

Implication: Due to this erratum, the processor may transfer control to an unintended address. Results are unpredictable, depending on the particular application, and can range from no effect to the unexpected termination of the application due to an exception. Intel has observed this erratum only in a focused testing environment. Intel has not observed any commercially available operating system, application, or compiler that makes use of or generates this instruction.

Workaround: If the other seven general purpose registers are unavailable for use, and it is necessary to do a CALL via the ESP register, first push ESP onto the stack, then perform an *indirect* call using ESP (e.g., CALL [ESP]). The saved version of ESP should be popped off the stack after the call returns.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B17. Built-in Self Test Always Gives Nonzero Result

Problem: The Built-in Self Test (BIST) of the Mobile Pentium II processor does not give a zero result to indicate a passing test. Regardless of pass or fail status, bit 6 of the BIST result in the EAX register after running BIST is set.

Implication: Software which relies on a zero result to indicate a passing BIST will indicate BIST failure.

Workaround: Mask bit 6 of the BIST result register when analyzing BIST results.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B18. Cache State Corruption in the Presence of Page A/D-bit Setting and Snoop Traffic

Problem: If an operating system uses the Page Access and/or Dirty bit feature implemented in the Intel architecture and there is a significant amount of snoop traffic on the bus, while the processor is setting the Access and/or Dirty bit the processor may inappropriately change a single L1 cache line to the modified state.

Implication: The occurrence of this erratum may result in cache incoherency, which may cause parity errors, data corruption (with no parity error), unexpected application or operating system termination, or system hangs.

Workaround: It is possible for BIOS code to contain a workaround for this erratum.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B19. Snoop Cycle Generates Spurious Machine Check Exception

Problem: The processor may incorrectly generate a Machine Check Exception (MCE) when it processes a snoop access that does not hit the L1 data cache. Due to an internal logic error, this type of snoop cycle may still check data parity on undriven data lines. The processor generates a spurious machine check exception as a result of this unnecessary parity check.

Implication: A spurious machine check exception may result in an unexpected system halt if Machine Check Exception reporting is enabled in the operating system.

Workaround: It is possible for BIOS code to contain a workaround for this erratum. This workaround would fix the erratum; however, the data parity error will still be reported.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B20. *MOVD/MOVQ Instruction Writes to Memory Prematurely*

Problem: When an instruction encounters a fault, the faulting instruction should not modify any CPU or system state. However, when the MMX™ technology store instructions MOVD and MOVQ encounter any of the following events, it is possible for the store to be committed to memory even though it should be canceled:

1. If CR0.EM = 1 (Emulation bit), then the store could happen prior to the triggered invalid opcode exception.
2. If the floating-point Top-of-Stack (FP TOS) is not zero, then the store could happen prior to executing the processor assist routine that sets the FP TOS to zero.
3. If there is an unmasked floating-point exception pending, then the store could happen prior to the triggered unmasked floating-point exception.
4. If CR0.TS = 1 (Task Switched bit), then the store could happen prior to the triggered Device Not Available (DNA) exception.

If the MOVD/MOVQ instruction is restarted after handling any of the above events, then the store will be performed again, overwriting with the expected data. The instruction will not be restarted after event 1. The instruction will definitely be restarted after events 2 and 4. The instruction may or may not be restarted after event 3, depending on the specific exception handler.

Implication: This erratum causes unpredictable behavior in an application if MOVD/MOVQ instructions are used to manipulate semaphores for multiprocessor synchronization, or if these MMX instructions are used to write to uncacheable memory or memory mapped I/O that has side effects, e.g., graphics devices. This erratum is completely transparent to all applications that do not have these characteristics. When each of the above conditions are analyzed:

1. Setting the CR0.EM bit forces all floating-point/MMX instructions to be handled by software emulation. The MOVD/MOVQ instruction, which is an MMX instruction, would be considered an invalid instruction. Operating systems typically terminates the application after getting the expected invalid opcode fault.
2. The FP TOS not equal to 0 case only occurs when the MOVD/MOVQ store is the first MMX instruction in an MMX technology routine and the previous floating-point routine did not clean up the floating-point states properly when it exited. Floating-point routines commonly leave TOS to 0 prior to exiting. For a store to be executed as the first MMX instruction in an MMX technology routine following a floating-point routine, the software would be implementing instruction level intermixing of floating-point and MMX instructions. Intel does not recommend this practice.
3. The unmasked floating-point exception case only occurs if the store is the first MMX technology instruction in an MMX technology routine and the previous floating-point routine exited with an unmasked floating-point exception pending. Again, for a store to be executed as the first MMX instruction in an MMX technology routine following a floating-point routine, the software would be implementing instruction level intermixing of floating-point and MMX instructions. Intel does not recommend this practice.

Device Not Available (DNA) exceptions occur naturally when a task switch is made between two tasks that use either floating-point instructions and/or MMX instructions. For this erratum, in the event of the DNA exception, data from the prior task may be temporarily stored to the present task's program state.

Workaround: Do not use MMX instructions to manipulate semaphores for multiprocessor synchronization. Do not use MOVD/MOVQ instructions to write directly to I/O devices if doing so triggers user visible side effects. An OS can prevent old data from being stored to a new task's program state by cleansing the FPU explicitly after every task switch. Follow Intel's recommended programming paradigms in the Intel Architecture Optimization Manual for writing MMX technology programs. Specifically, do not mix floating-point and MMX instructions. When transitioning to new a MMX technology routine, begin with an instruction that does not depend on the prior state of either the MMX technology registers or the floating-point registers, such as a load or PXOR mm0, mm0. Be sure that the FP TOS is clear before using MMX instructions.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B21. Memory Type Undefined for Nonmemory Operations

Problem: The Memory Type field for nonmemory transactions such as I/O and Special Cycles are undefined. Although the Memory Type attribute for nonmemory operations logically should (and usually does) manifest itself as UC, this feature is not designed into the implementation and is therefore inconsistent.

Implication: Bus agents may decode a non-UC memory type for nonmemory bus transactions.

Workaround: Bus agents must consider transaction type to determine the validity of the Memory Type field for a transaction.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B22. FP Data Operand Pointer May Not Be Zero After Power On or Reset

Problem: The FP Data Operand Pointer, as specified, should be reset to zero upon power on or Reset by the processor. Due to this erratum, the FP Data Operand Pointer may be nonzero after power on or Reset.

Implication: Software which uses the FP Data Operand Pointer and count on its value being zero after power on or Reset without first executing an FINIT/FNINIT instruction will use an incorrect value, resulting on incorrect behavior of the software.

Workaround: Software should follow the recommendation in Section 8.2 of the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide* (Order Number 243192). This recommendation states that if the FPU will be used, software-initialization code should execute an FINIT/FNINIT instruction following a hardware reset. This will correctly clear the FP Data Operand Pointer to zero.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B23. *MOVD Following Zeroing Instruction Can Cause Incorrect Result*

Problem: An incorrect result may be calculated after the following circumstances occur:

1. A register has been zeroed with either a SUB reg, reg instruction or an XOR reg, reg instruction,
2. A value is moved with sign extension into the same register's lower 16 bits; or a signed integer multiply is performed to the same register's lower 16 bits,
3. This register is then copied to an MMX™ technology register using the MOVD instruction prior to any other operations on the sign-extended value.

Specifically, the sign may be incorrectly extended into bits 16-31 of the MMX technology register. Only the MMX technology register is affected by this erratum.

The erratum only occurs when the 3 following steps occur in the order shown. The erratum may occur with up to 40 intervening instructions that do not modify the sign-extended value between steps 2 and 3.

1. XOR EAX, EAX
or SUB EAX, EAX
2. MOVSX AX, BL
or MOVSX AX, byte ptr <memory address> or MOVSX AX, BX
or MOVSX AX, word ptr <memory address> or IMUL BL (AX implicit, opcode F6 /5)
or IMUL byte ptr <memory address> (AX implicit, opcode F6 /5) or IMUL AX, BX (opcode 0F AF /r)
or IMUL AX, word ptr <memory address> (opcode 0F AF /r) or IMUL AX, BX, 16 (opcode 6B /r ib)
or IMUL AX, word ptr <memory address>, 16 (opcode 6B /r ib) or IMUL AX, 8 (opcode 6B /r ib)
or IMUL AX, BX, 1024 (opcode 69 /r iw)
or IMUL AX, word ptr <memory address>, 1024 (opcode 69 /r iw) or IMUL AX, 1024 (opcode 69 /r iw)
or CBW
3. MOVD MM0, EAX

Note that the values for immediate byte/words are merely representative (i.e., 8, 16, 1024) and that any value in the range for the size may be affected. Also, note that this erratum may occur with "EAX" replaced with any 32-bit general purpose register, and "AX" with the corresponding 16-bit version of that replacement. "BL" or "BX" can be replaced with any 8-bit or 16-bit general purpose register. The CBW and IMUL (opcode F6 /5) instructions are specific to the EAX register only.

In the example, EAX is forced to contain 0 by the XOR or SUB instructions. Since the MOVSX, IMUL and CBW instructions listed should modify only bits 15:8 of EAX by sign extension, bits 31:16 of EAX should always contain 0. This implies that when MOVD copies EAX to MM0, bits 31:16 of MM0 should also be 0. Under certain scenarios, bits 31:16 of MM0 are not 0, but are replicas of bit 15 (the 16th bit) of AX. This is noticeable when the value in AX after the MOVSX, IMUL or CBW instruction is negative, i.e., bit 15 of AX is a 1.

When AX is positive (bit 15 of AX is a 0), MOVD will always produce the correct answer. If AX is negative (bit 15 of AX is a 1), MOVD may produce the right answer or the wrong answer depending on the point in time when the MOVD instruction is executed in relation to the MOVSX, IMUL or CBW instruction.

Implication: The effect of incorrect execution will vary from unnoticeable, due to the code sequence discarding the incorrect bits, to an application failure. If the MMX technology-enabled application in which MOVD is used to manipulate pixels, it is possible for one or more pixels to exhibit the wrong color or position momentarily. It is also possible for a computational application that uses the MOVD instruction in the manner described above to produce incorrect data. Note that this data may cause an unexpected page fault or general protection fault.

Workaround: There are two possible workarounds for this erratum:

1. Rather than using the MOVXS-MOVD or CBW-MOVD pairing to handle one variable at a time, use the sign extension capabilities (PSRAW, etc.) within MMX technology for operating on multiple variables. This would result in higher performance as well.
2. Insert another operation that modifies or copies the sign-extended value between the MOVXS/IMUL/CBW instruction and the MOVD instruction as in the example below:

```
XOR EAX, EAX (or SUB EAX, EAX)
MOVXS AX, BL (or other MOVXS, other IMUL or CBW instruction)
*MOV EAX, EAX
MOVD MM0, EAX
```

*Note: MOV EAX, EAX is used here as it is fairly generic. Again, EAX can be any 32-bit register.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B24. Premature Execution of a Load Operation Prior to Exception Handler Invocation

Problem: This erratum can occur with any of the following situations:

1. If an instruction that performs a memory load causes a code segment limit violation
2. If a waiting floating-point instruction or MMX™ instruction that performs a memory load has a floating-point exception pending
3. If an MMX instruction that performs a memory load and has either CR0.EM =1 (Emulation bit set), or a floating-point Top-of-Stack (FP TOS) not equal to 0, or a DNA exception pending

If any of the above circumstances, occur it is possible that the load portion of the instruction will have executed before the exception handler is entered.

Implication: In normal code execution where the target of the load operation is to write back memory there is no impact from the load being prematurely executed, nor from the restart and subsequent re-execution of that instruction by the exception handler. If the target of the load is to uncached memory that has a system side-effect, restarting the instruction may cause unexpected system behavior due to the repetition of the side-effect.

Workaround: Code which performs loads from memory that has side-effects can effectively workaround this behavior by using simple integer-based load instructions when accessing side-effect memory and by ensuring that all code is written such that a code segment limit violation cannot occur as a part of reading from side-effect memory.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B25. Read Portion of RMW Instruction May Execute Twice

Problem: When the Mobile Pentium II processor executes a read-modify-write (RMW) arithmetic instruction, with memory as the destination, it is possible for a page fault to occur during the execution of the store on the memory operand after the read operation has completed but before the write operation completes.

If the memory targeted for the instruction is UC (uncached), memory will observe the occurrence of the initial load before the page fault handler and again if the instruction is restarted.

Implication: This erratum has no effect if the memory targeted for the RMW instruction has no side-effects. If, however, the load targets a memory region that has side-effects, multiple occurrences of the initial load may lead to unpredictable system behavior.

Workaround: Hardware and software developers who write device drivers for custom hardware that may have a side-effect style of design should use simple loads and simple stores to transfer data to and from the device. Then, the memory location will simply be read twice with no additional implications.

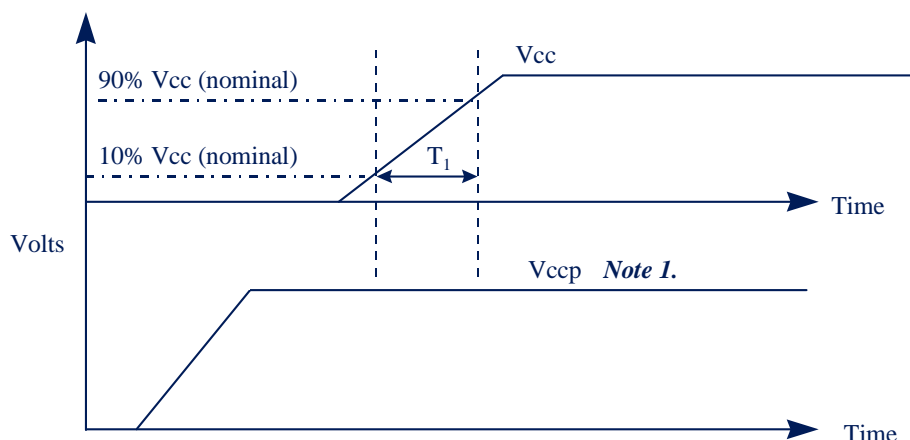
Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B26. Voltage Sequencing During Power Up

Problem: The processor uses the PWRGOOD signal to ensure that no voltage sequencing issues arise; no pin assertions should cause the processor to change its behavior until this signal is asserted, when all power supplies and clocks to the processor are valid and stable. However, if V_{CC} ramps up before V_{CCP} , it will cause the processor to enter an invalid state.

Implication: If this erratum occurs, the system may boot normally however, the L2 cache may not be initialized.

Workaround: For Mobile Pentium II processors, ensure that V_{CC} ramps after or with the L2 cache IO voltage (V_{CCP}). If V_{CCP} ramps after the V_{CC} , a delay circuit can be placed on the core V_{CC} to ensure that V_{CC} ramps at the same time or after V_{CCP} as shown below in Figure 1.



Symbol	Parameter	Min	Max	Notes
T ₁	VCC Ramp Time	200 uSec		Measured between 10% VCC (nom) to 90% VCC (nom)
Note 1: VCCP must be greater than VCC during the interval in which VCC rises from 20% VCC to 80% VCC.				

Figure 1. VCC/VCCP Timing Requirements (Mobile)

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B27. Intervening Writeback May Occur During Locked Transaction

Problem: During a transaction which has the LOCK# signal asserted (e.g., a locked transaction), there is a potential for an explicit writeback caused by a previous transaction to complete while the bus is locked. The explicit writeback will only be issued by the processor which has locked the bus, and the lock signal will not be deasserted until the locked transaction completes, but the atomicity of a lock may be compromised by this erratum. Note that the explicit writeback is an expected cycle, and no memory ordering violations will occur. This erratum is, however, a violation of the bus lock protocol.

Implication: A chipset or third-party agent (TPA) which tracks bus transactions in such a way that locked transactions may only consist of a read-write or read-read-write-write locked sequence, with no transactions intervening, may lose synchronization of state due to the intervening explicit writeback. Systems using chipsets or TPAs which can accept the intervening transaction will not be affected.

Workaround: The bus tracking logic of all devices on the system bus should allow for the occurrence of an intervening transaction during a locked transaction.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B28. MC2_STATUS MSR Has Model-Specific Error Code and Machine Check Architecture Error Code Reversed

Problem: The *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, documents that for the MC1_STATUS MSR, bits 15:0 contain the MCA (machine-check architecture) error code field and bits 31:16 contain the model-specific error code field. However, for the MC2_STATUS MSR, these bits have been reversed. For the MC2_STATUS MSR, bits 15:0 contain the model-specific error code field and bits 31:16 contain the MCA error code field.

Implication: A machine check error may be decoded incorrectly if this erratum on the MC2_STATUS MSR is not taken into account.

Workaround: When decoding the MC2_STATUS MSR, reverse the two error fields.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B29. Mixed Cacheability of Lock Variables Is Problematic in MP Systems

Problem: This errata only affects multiprocessor systems where a lock variable address is marked cacheable in one processor and uncacheable in any others. The processors which have it marked uncacheable may stall indefinitely when accessing the lock variable. The stall is only encountered if:

- One processor has the lock variable cached, and is attempting to execute a cache lock.
- The processor which has that address cached has it cached in its L2 only.
- Other processors, meanwhile, issue back to back accesses to that same address on the bus.

Implication: MP systems where all processors either use cache locks or consistent locks to uncacheable space will not encounter this problem. If, however, a lock variable's cacheability varies in different processors, and several processors are all attempting to perform the lock simultaneously, an indefinite stall may be experienced by the processors which have it marked uncacheable in locking the variable (if the conditions above are satisfied). Intel has only encountered this problem in focus testing with artificially generated external events. Intel has not currently identified any commercial software which exhibits this problem.

Workaround: Follow a homogenous model for the memory type range registers (MTRRs), ensuring that all processors have the same cacheability attributes for each region of memory; do not use locks whose memory type is cacheable on one processor, and uncacheable on others. Avoid page table aliasing, which may produce a nonhomogenous memory model.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B30. Thermal Sensor May Assert SMBALERT# Incorrectly

Problem: The Mobile Pentium II processor and the Pentium II Processor Mobile Module have a thermal sensor that monitors the processor core's temperature. Please note that desktop systems could have a similar thermal device. The thermal sensor asserts SMBALERT# if the processor temperature exceeds the temperature limits set in the Alarm Threshold Registers (T_{HIGH} , T_{LOW}). It also sets the corresponding Status Register bits to identify the cause of the interrupt. Figure 1 gives one example of the how the SMBALERT# signal could be used in a system.

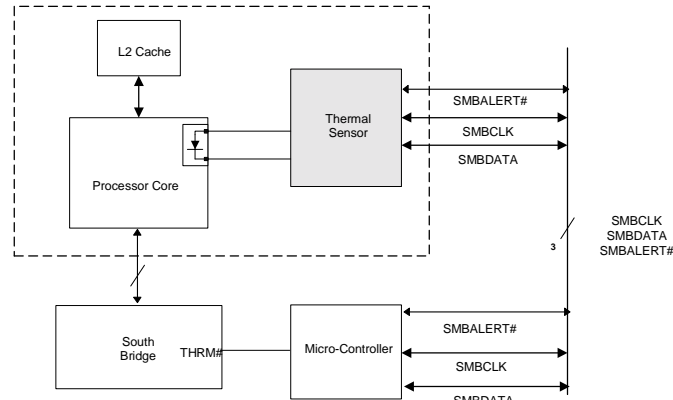


Figure 1. An Example of Microcontroller Driven Thermal Management

Under the conditions described below, the thermal sensor incorrectly generates the SMBALERT# interrupt. All of the following conditions must be met to trigger a false interrupt:

1. The thermal sensor must be in auto-convert mode.
2. The absolute value of the difference between the current temperature reading and the T_{HIGH} or T_{LOW} limit value must be less than or equal to 8° C.
3. The current temperature reading must be different from the previous reading.

With a false assertion of SMBALERT#, the corresponding bit in the Status Register (L_{HIGH} , L_{LOW} , R_{HIGH} , and R_{LOW}) also will be incorrect.

Implication: There is no system impact from this erratum if temperature polling is used for processor thermal management. If the SMBALERT# interrupt is employed to manage processor thermal sensing, then servicing the false interrupt may result in premature system action depending on the software and hardware implementations used. The rate of the false interrupts is less than the auto-convert rate of the thermal sensor.

Workaround: Three different (mutually exclusive) workarounds are possible:

1. Before servicing an interrupt from the thermal sensor, read and compare the processor thermal reading with the threshold limits (T_{HIGH} or T_{LOW}). Figures 2 and 3 provide basic flowcharts for the implementation of this workaround in an interrupt driven system.
2. If the firmware implemented polls the Status Register only, then before taking any action, re-read the temperature register and do a comparison with the alarm threshold limits (T_{HIGH} or T_{LOW}) to determine if the value is actually still within the temperature window.

3. Use a temperature polling scheme to monitor the processor temperature.

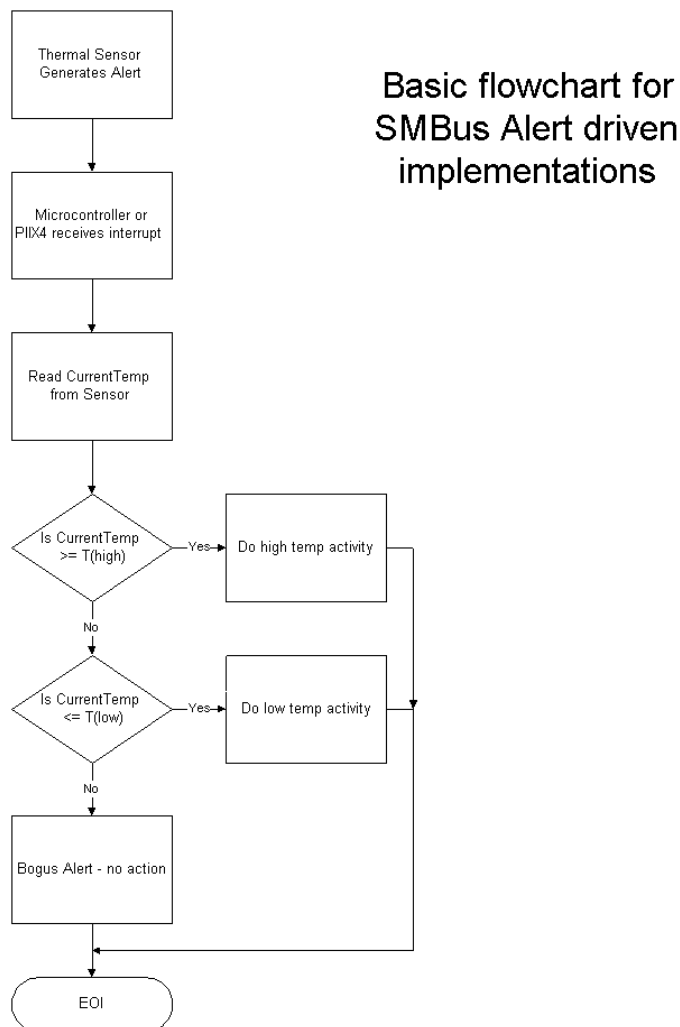


Figure 2. Workaround Flowchart: SMBALERT#-Driven System

Basic flowchart for system interrupt driven implementations

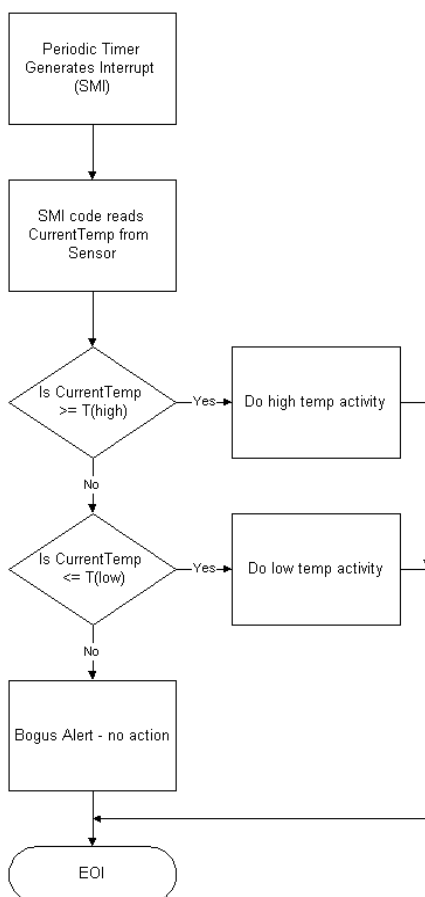


Figure 3. Workaround Flowchart: SMI#-Driven System

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B31. *MOV With Debug Register Causes Debug Exception*

Problem: When in V86 mode, if a MOV instruction is executed on debug registers, a general-protection exception (#GP) should be generated, as documented in the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, Section 15.2. However, in the case when the general detect enable flag (GD) bit is set, the observed behavior is that a debug exception (#DB) is generated instead.

Implication: With debug-register protection enabled (e.g., the GD bit set), when attempting to execute a MOV on debug registers in V86 mode, a debug exception will be generated instead of the expected general-protection fault.

Workaround: In general, operating systems do not set the GD bit when they are in V86 mode. The GD bit is generally set and used by debuggers. The debug exception handler should check that the exception did not occur in V86 mode before continuing. If the exception did occur in V86 mode, the exception may be directed to the general-protection exception handler.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B32. *Upper Four PAT Entries Not Usable With Mode B or Mode C Paging*

Problem: The Page Attribute Table (PAT) contains eight entries, which must all be initialized and considered when setting up memory types for the Mobile Pentium II processor. However, in Mode B or Mode C paging, the upper four entries do not function correctly for 4-Kbyte pages. Specifically, bit seven of page table entries that translate addresses to 4-Kbyte pages should be used as the upper bit of a three-bit index to determine the PAT entry that specifies the memory type for the page. When Mode B (CR4.PSE = 1) and/or Mode C (CR4.PAE) are enabled, the processor forces this bit to zero when determining the memory type regardless of the value in the page table entry. The upper four entries of the PAT function correctly for 2-Mbyte and 4-Mbyte large pages (specified by bit 12 of the page directory entry for those translations).

Implication: Only the lower four PAT entries are useful for 4KB translations when Mode B or C paging is used. In Mode A paging (4-Kbyte pages only), all eight entries may be used. All eight entries may be used for large pages in Mode B or C paging.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B33. UC Write May be Reordered Around a Cacheable Write

Problem: After a write occurs to a UC (uncacheable) region of memory, there exists a small window of opportunity where a subsequent write transaction targeted for a UC memory region may be reordered in front of a write targeted to a region of cacheable memory. This erratum can only occur during the following sequence of bus transactions:

1. A write to memory mapped as UC occurs
2. A write to memory mapped as cacheable (WB or WT) which is present in Shared or Invalid state in the L2 cache occurs
3. During the bus snoop of the cacheable line, another store to UC memory occurs

Implication: If this erratum occurs, the second UC write will be observed on the bus prior to the Bus Invalidate Line (BIL) or Bus Read Invalidate Line (BRIL) transaction for the cacheable write. This presents a small window of opportunity for a fast bus-mastering I/O device which triggers an action based on the second UC write to arbitrate and gain ownership of the bus prior to the completion of the cacheable write, possibly retrieving stale data.

Workaround: It is possible for BIOS code to contain a workaround for this erratum.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B34. Incorrect Memory Type May be Used When MTRRs Are Disabled

Problem: If the Memory Type Range Registers (MTRRs) are disabled without setting the CR0.CD bit to disable caching, and the Page Attribute Table (PAT) entries are left in their default setting, which includes UC- memory type (PCD = 1, PWT = 0; see the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, for details), data for entries set to UC- will be cached as if the memory type were writeback (WB). Also, if the page tables are set to a memory type other than UC-, then the effective memory type used will be that specified by the page tables and PAT. Any regions of memory normally forced to UC by the MTRRs (such as the VGA video region) may now be incorrectly cached and speculatively accessed.

Even if the CR0.CD bit is correctly set when the MTRRs are disabled and the PAT is left in its default state, then retries and out of order retirement of UC accesses may occur, contrary to the strong ordering expected for these transactions.

Implication: The occurrence of this erratum may result in the use of incorrect data and unpredictable processor behavior when running with the MTRRs disabled. Interaction between the mouse, cursor, and VGA video display leading to video corruption may occur as a symptom of this erratum as well.

Workaround: Ensure that when the MTRRs are disabled, the CR0.CD bit is set to disable caching. This recommendation is described in the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*. If it is necessary to disable the MTRRs, first clear the PAT register before setting the CR0.CD bit, flushing the caches, and disabling the MTRRs to ensure that UC memory type is always returned and strong ordering is maintained.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B35. Misprediction in Program Flow May Cause Unexpected Instruction Execution

Problem: To optimize performance through dynamic execution technology, the P6 architecture has the ability to predict program flow. In the event of a misprediction, the processor will normally clear the incorrect prediction, adjust the EIP to the correct location, and flush out any instructions it may have fetched from the misprediction. In circumstances where a branch misprediction occurs, the correct target of the branch has already been opportunistically fetched into the streaming buffers, and the L2 cycle caused by the evicted cache line is retried by the L2 cache, the processor may fail to flush out the retirement unit before the speculative program flow is committed to a permanent state.

Implication: The results of this erratum may range from no effect to unpredictable application or OS failure. Manifestations of this failure may result in:

- Unexpected values in EIP
- Faults or traps (e.g., page faults) on instructions that do not normally cause faults
- Faults in the middle of instructions
- Unexplained values in registers/memory at the correct EIP

Workaround: It is possible for BIOS code to contain a workaround for this erratum.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B36. Data Breakpoint Exception in a Displacement Relative Near Call May Corrupt EIP

Problem: If a misaligned data breakpoint is programmed to the same cache line as the memory location where the stack push of a near call is performed and any data breakpoints are enabled, the processor will update the stack and ESP appropriately, but may skip the code at the destination of the call. Hence, program execution will continue with the next instruction immediately following the call, instead of the target of the call.

Implication: The failure mechanism for this erratum is that the call would not be taken; therefore, instructions in the called subroutine would not be executed. As a result, any code relying on the execution of the subroutine will behave unpredictably.

Workaround: Whether enabled or not, do not program a misaligned data breakpoint to the same cache line on the stack where the push for the near call is performed.

Status: For the stepping affected see the *Summary of Changes* at the beginning of this section.

B37. System Bus ECC Not Functional With 2:1 Ratio

Problem: If a processor is underclocked at a core frequency to system bus frequency ratio of 2:1 and system bus ECC is enabled, the system bus ECC detection and correction will negatively affect internal timing dependencies.

Implication: If system bus ECC is enabled, and the processor is underclocked at a 2:1 ratio, the system may behave unpredictably due to these timing dependencies.

Workaround: All bus agents that support system bus ECC must disable it when a 2:1 ratio is used.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B38. Fault on REP CMPS/SCAS Operation May Cause Incorrect EIP

Problem: If either a General Protection Fault, Alignment Check Fault or Machine Check Exception occur during the first iteration of a REP CMPS or a REP SCAS instruction, an incorrect EIP may be pushed onto the stack of the event handler if all the following conditions are true:

- The event occurs on the initial load performed by the instruction(s)
- The condition of the zero flag before the repeat instruction happens to be opposite of the repeat condition (e.g., REP/REPE/REPZ CMPS/SCAS with ZF = 0 or RENE/REPZ CMPS/SCAS with ZF = 1)
- The faulting micro-op and a particular micro-op of the REP instruction are retired in the retirement unit in a specific sequence

The EIP will point to the instruction following the REP CMPS/SCAS instead of pointing to the faulting instruction.

Implication: The result of the incorrect EIP may range from no effect to unexpected application/OS behavior.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B39. RDMSR or WRMSR to Invalid MSR Address May Not Cause GP Fault

Problem: The RDMSR and WRMSR instructions allow reading or writing of MSRs (Model Specific Registers) based on the index number placed in ECX. The processor should reject access to any reserved or unimplemented MSRs by generating #GP(0). However, there are some invalid MSR addresses for which the processor will not generate #GP(0).

Implication: For RDMSR, undefined values will be read into EDX:EAX. For WRMSR, undefined processor behavior may result.

Workaround: Do not use invalid MSR addresses with RDMSR or WRMSR.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B40. SYSENTER/SYSEXIT Instructions Can Implicitly Load “Null Segment Selector” to SS and CS Registers

Problem: According to the processor specification, attempting to load a null segment selector into the CS and SS segment registers should generate a General Protection Fault (#GP). Although loading a null segment selector to the other segment registers is allowed, the processor will generate an exception when the segment register holding a null selector is used to access memory.

However, the SYSENTER instruction can implicitly load a null value to the SS segment selector. This can occur if the value in SYSENTER_CS_MSR is between FFF8h and FFFBh when the SYSENTER instruction is executed. This behavior is part of the SYSENTER/SYSEXIT instruction definition; the content of the SYSTEM_CS_MSR is always incremented by 8 before it is loaded into the SS. This operation will set the null bit in the segment selector if a null result is generated, but it does not generate a #GP on the SYSENTER instruction itself. An exception will be generated as expected when the SS register is used to access memory, however.

The SYSEXIT instruction will also exhibit this behavior for both CS and SS when executed with the value in SYSENTER_CS_MSR between FFF0h and FFF3h, or between FFE8h and FFEb, inclusive.

Implication: These instructions are intended for operating system use. If this erratum occurs (and the OS does not ensure that the processor never has a null segment selector in the SS or CS segment registers), the processor's behavior may become unpredictable, possibly resulting in system failure.

Workaround: Do not initialize the SYSTEM_CS_MSR with the values between FFF8h and FFFBh, FFF0h and FFF3h, or FFE8h and FFEb before executing SYSENTER or SYSEXIT.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B41. PRELOAD Followed by EXTEST Does Not Load Boundary Scan Data

Problem: According to the IEEE 1149.1 Standard, the EXTEST instruction would use data “typically loaded onto the latched parallel outputs of boundary-scan shift-register stages using the SAMPLE/PRELOAD instruction prior to the selection of the EXTEST instruction.” As a result of this erratum, this method cannot be used to load the data onto the outputs.

Implication: Using the PRELOAD instruction prior to the EXTEST instruction will not produce expected data after the completion of EXTEST.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B42. Far Jump to New TSS With D-bit Cleared May Cause System Hang

Problem: A task switch may be performed by executing a far jump through a task gate or to a new Task State Segment (TSS) directly. Normally, when such a jump to a new TSS occurs, the D-bit (which indicates that the page referenced by a Page Table Entry (PTE) has been modified) for the PTE which maps the location of the previous TSS will already be set and the processor will operate as expected. However, if the D-bit is clear at the time of the jump to the new TSS, the processor will hang.

Implication: If an OS is used which can clear the D-bit for system pages, and which jumps to a new TSS on a task switch, then a condition may occur which results in a system hang. Intel has not identified any commercial software which may encounter this condition; this erratum was discovered in a focused testing environment.

Workaround: Ensure that OS code does not clear the D-bit for system pages (including any pages that contain a task gate or TSS). Use task gates rather than jumping to a new TSS when performing a task switch.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B43. Incorrect Chunk Ordering May Prevent Execution of the Machine Check Exception Handler After BINIT#

Problem: If a catastrophic bus error is detected which results in a BINIT# assertion, and the BINIT# assertion is propagated to the processor's L2 cache at the same time that data is being sent to the processor, then the data may become corrupted in the processor's L1 cache.

Implication: Since BINIT# assertion is due to a catastrophic event on the bus, the corrupted data will not be used. However, it may prevent the processor from executing the Machine Check Exception (MCE) handler, causing the system to hang.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B44. Resume Flag May Not Be Cleared After Debug Exception

Problem: The Resume Flag (RF) is normally cleared by the processor after executing an instruction which causes a debug exception (#DB). In the process of determining whether the RF needs to be cleared after executing the instruction, the processor uses an internal register containing stale data. The stale data may unpredictably prevent the processor from clearing the RF.

Implication: If this erratum occurs, further debug exceptions will be disabled.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B45. Thermal Sensor Leakage Current May Exceed Specification

Problem: The thermal sensor contained on the Pentium II Processor Mobile Module violates the input leakage current specification of $\pm 1\mu\text{A}$ as specified in the SMBus Specification.

Implication: The thermal sensor incorporates input protection diodes on the SMBCLK and SMBDATA signals for ESD protection. The protection diodes can potentially clamp these lines to $\sim 0.6\text{ V}$ when the V₃ voltage supply to the module is powered off. Hence, when module is powered down, the thermal sensor on the Pentium II Processor Mobile Module may prevent SMBus transactions from occurring. When SMBus devices are powered from different voltage power planes, they are typically isolated from one another. This erratum will not affect designs using this isolation technique.

Workaround: It is recommended to have SMBus devices that are powered by a source other than V₃, be isolated from the Pentium II Processor Mobile Module SMBus to ensure that this erratum does not occur.

Status: For the steppings affected see the Summary Table of Changes at the beginning of this section.

B46. Processor May Return Invalid Parameters on Execution of the CPUID Instruction

Problem: The Mobile Module based Pentium II processor with on-die L2 cache may return an incorrect "Maximum CPUID Input Value." The Mobile Module based Pentium II processor with on-die L2 cache is specified to return a value of 2 in the EAX register when the CPUID instruction is executed with EAX=0; however, this erratum may result in the value of 3 being returned in EAX. It is also possible that bit 18 of the EDX register will be set to a 1 when CPUID is executed with EAX=1. This bit is defined as reserved for the Mobile Module based Pentium II processor with on-die L2 cache, but is expected to be set to zero. If CPUID were to be executed on the Mobile Module based Pentium II processor with on-die L2 cache with EAX=3 the processor should return the cache parameters in the integer registers (EAX, EBX, ECX, EDX); however, the processors affected by this erratum will return undefined values in the integer registers.

Implication: Intel has not seen any software failures as a result of this erratum; however, since software written for the Mobile Module based Pentium II processor with on-die L2 cache will not be expecting to see a Maximum CPUID Input Value greater than 2, it is not possible to predict how software will behave on processors with this erratum. Software using the fact that bit 18 in the feature flags is set, to determine the presence of the Pentium® III processor serial number feature, without also verifying that it is executing on a Pentium III processor, may incorrectly believe that the Mobile Module based Pentium II processor with on-die L2 cache, support the processor serial number feature. However, the values the CPUID instruction returns when CPUID is executed with EAX=3 will not be the processor serial number and will be undefined.

Workaround: It is possible for BIOS code to contain a workaround for this erratum.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B47. Internal Cache Protocol Violation May Cause System Hang

Problem: A Mobile Pentium II processor based system may hang due to an internal cache protocol violation. During multiple transactions targeted at the same cacheline, there exists a small window of time such that the processor's internal timings align to create a livelock situation. The scenario, which results in the erratum, is summarized below:

Scenario:

1. A snoopable transaction is issued to address A. This snoopable transaction can be issued by the processor or the chipset.
2. The snoopable transaction hits a modified line in the processor's L1 data cache.
3. The processor issues two code fetches from the L2 cache before the snoopable transaction reaches the top of the In-Order Queue and before the snoopable transaction's modified L1 cache line containing address A is brought out on the system bus.

At the same time, a locked access to the L1 cache occurs.

Implication: A Mobile Pentium II processor may cause a system to hang if the above listed sequence of events occur. The probability of encountering this erratum increases with I/O queue depth greater than four.

Workaround: It is possible for BIOS code to contain a workaround for this erratum.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B48. GP# Fault on WSMSR to ROB_CR_BKUPTMPDR6

Problem: Writing a '1' to unimplemented bit(s) in the ROB_CR_BKUPTMPDR6 MSR (offset 1E0H) will result in a general protection fault (GP#).

Implication: The normal process used to write an MSR is to read the MSR using RDMSR, modify the bit(s) of interest, and then to write the MSR using WRMSR. Because of this erratum, this process may result in a GP# fault when used to modify the ROB_CR_BKUPTMPDR6 MSR

Workaround: When writing to ROB_CR_BKUPTMPDR6 all unimplemented bits must be '0.' Implemented bits may be set as '0' or '1' as desired.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B49. Machine Check Exception May Occur Due to Improper Line Eviction in the IFU

Problem: The Mobile Pentium II processor is designed to signal an unrecoverable Machine Check Exception (MCE) as a consistency checking mechanism. Under a complex set of circumstances involving multiple speculative branches and memory accesses there exists a one cycle long window in which the processor may signal a MCE in the Instruction Fetch Unit (IFU) because instructions previously decoded have been evicted from the IFU. The one cycle long window is opened when an opportunistic fetch receives a partial hit on a previously executed but not as yet completed store resident in the store buffer. The resulting partial hit erroneously causes the eviction of a line from the IFU at a time when the processor is expecting the line to still be present. If the MCE for this particular IFU event is disabled, execution will continue normally.

Implication: While this erratum may occur on a system with any number of processors, the probability of occurrence increases with the number of processors. If this erratum does occur, a machine check exception will result. Note systems that implement an operating system that does not enable the Machine Check Architecture will be completely unaffected by this erratum, e.g., Windows® 95 and Windows 98.

Workaround: It is possible for BIOS code to contain a workaround for this erratum.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B50. System May Hang When the Mobile Pentium® II Processor Is in the Deep Sleep State

Problem: Under normal operating conditions, when the Mobile Pentium II processor is in the Quick Start state, stopping the BCLK input puts the processor into the Deep Sleep state. For proper operation, the processor should return to the Quick Start state when the BCLK input is re-started. However, the processor may not exit the Deep Sleep state if the BCLK reassertion occurs less than 8 µS after the previous BCLK deassertion.

Implication: For a system BIOS that enables the Mobile Pentium II processor Deep Sleep state, the system may hang. Intel has not seen the system hang due to this erratum if the BCLK reassertion occurs greater than 8µS after the BCLK deassertion.

Workaround: It is possible to implement workarounds to avoid this erratum by limiting the processor BCLK deassertion. The following are three examples:

1. Disable the C3 state via ACPI BIOS modification to restrict the processor for entering the Deep Sleep state. A similar approach can be implemented in an APM BIOS.
2. If a system is designed with an Intel PXII4 bridge, register bit "STPCLK_EN" (I/O address = Base + (10h), bit 12. Please refer to the *82371AB PCI-TO-ISA/IDE Xcelerator (PIIX4)* datasheet for more details. It can be set to "0" via a BIOS modification to prevent the processor BCLK from being deasserted in the C3 state. The S1/POS state is not affected by the STPCLK_EN bit and the use of Deep Sleep in the PIIX4 implementation of the S1/POS state does not trigger this erratum. Hence, the processor is kept out of the Deep Sleep state when it is vulnerable to this erratum. For other systems using a different PCI to ISA Bridge, a similar approach can be implemented if such a bit is available.

It is possible to add a hardware delay on the BCLK reassertion when the CPU is in the Deep Sleep state. For example, if the assertion of BCLK is controlled by the PIIX4 CPU_STP#pin, sufficient delay can be added to the CPU_STP# rising edge to delay only the BCLK reassertion which prevent the system hang. Similar approaches can be implemented for other systems with different configurations.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B51. VTOL Pin May Be Driven to Ground on the Mobile Pentium® II Processor

Problem: The VTOL pin could be used in system designs to distinguish a 400 MHz processor from a processor at 366 MHz or below. This pin is specified to be 1.6 V tolerant. In normal operating conditions, this pin is at high impedance for a processor at 366 MHz and below. When pulled up to 1.6 V, the output voltage of this pin is a logical high level at about 1.6 V. However, the CPU may drive this pin to ground.

Implication: In a system that uses this pin to control the VR logic, the VR logic may misinterpret the processor to be a 400-MHz CPU, therefore fail to set the correct processor core and other voltages. The behavior of the system under the incorrect voltages is unknown.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B52. Lower Bits of SMRAM SMBASE Register Cannot Be Written With an ITP

Problem: The System Management Base (SMBASE) register (7EF8H) stores the starting address of the System Management RAM (SMRAM). This register is used by the processor when it is in System Management Mode (SMM), and its contents serve as the memory base for code execution and data storage. The 32-bit SMBASE register can normally be programmed to any value. When programmed with an In-Target Probe (ITP), however, any attempt to set the lower 11 bits of SMBASE to anything other than zeros via the WRMSR instruction will cause the attempted write to fail.

Implication: When set via the ITP, any attempt to relocate SMRAM space must be made with 2 Kbyte alignment.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B53. Task Switch May Cause Wrong PTE and PDE Access Bit to be Set

Problem: If an operating system executes a task switch via a Task State Segment (TSS), and the TSS is wholly or partially located within a clean page (A and D bits clear) and the GDT entry for the new TSS is either misaligned across a cache line boundary or is in a clean page, the accessed and dirty bits for an incorrect page table/directory entry may be set.

Implication: An operating system that uses hardware task switching (or hardware task management) may encounter this erratum. The effect of the erratum depends on the alignment of the TSS and ranges from no anomalous behavior to unexpected errors.

Workaround: The operating system could align all TSSs to be within page boundaries and set the A and D bits for those pages to avoid this erratum. The operating system may alternately use software task management.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B54. Unsynchronized Cross-Modifying Code Operations Can Cause Unexpected Instruction Execution Results

Problem: The act of one processor, or system bus master, writing data into a currently executing code segment of a second processor with the intent of having the second processor execute that data as code is called cross-modifying code (XMC). XMC that does not force the second processor to execute a synchronizing instruction, prior to execution of the new code, is called unsynchronized XMC.

Software using unsynchronized XMC to modify the instruction byte stream of a processor can see unexpected instruction execution from the processor which is executing the modified code.

Implication: In this case, the phrase "unexpected execution behavior" encompasses the generation of most of the exceptions listed in the *Intel Architecture Software Developer's Manual Volume 3: System Programming Guide* including a General Protection Fault (GPF). In the event of a GPF the application executing the unsynchronized XMC operation would be terminated by the operating system.

Workaround: In order to avoid this erratum, programmers should use the XMC synchronization algorithm as detailed in the *Intel Architecture Software Developer's Manual Volume 3: System Programming Guide*, Section 7.1.3.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B55. *Deadlock May Occur Due To Illegal-Instruction/Page-Miss Combination*

Problem: Intel's 32-bit Instruction Set Architecture (ISA) utilizes most of the available op-code space; however some byte combinations remain undefined and are considered illegal instructions. Intel processors detect the attempted execution of illegal instructions and signal an exception. This exception is handled by the operating system and/or application software.

Under a complex set of internal and external conditions involving illegal instructions, a deadlock may occur within the processor. The necessary conditions for the deadlock involve:

1. The illegal instruction is executed.
2. Two page table walks occur within a narrow timing window coincident with the illegal instruction.

Implication: The illegal instructions involved in this erratum are unusual and invalid byte combinations that are not useful to application software or operating systems. These combinations are not normally generated in the course of software programming, nor are such sequences known by Intel to be generated in commercially available software and tools. Development tools (compilers, assemblers) do not generate this type of code sequence, and will normally flag such a sequence as an error. If this erratum occurs, the processor deadlock condition will occur and result in a system hang. Code execution cannot continue without a system RESET.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B56. *FLUSH# Assertion Following STPCLK# May Prevent CPU Clocks From Stopping*

Problem: If FLUSH# is asserted after STPCLK# is asserted, the cache flush operation will not occur until after STPCLK# is de-asserted. Furthermore, the pending flush will prevent the processor from entering the Sleep state, since the flush operation must complete prior to the processor entering the Sleep state.

Implication: Following SLP# assertion, processor power dissipation may be higher than expected. Furthermore, if the source to the processor's input bus clock (BCLK) is removed, normally resulting in a transition to the Deep Sleep state, the processor may shutdown improperly. The ensuing attempt to wake up the processor will result in unpredictable behavior and may cause the system to hang.

Workaround: For systems that use the FLUSH# input signal and Deep Sleep state of the processor, ensure that FLUSH# is not asserted while STPCLK# is asserted.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B57. Floating-Point Exception Condition Can Be Deferred

Problem: A floating-point instruction that causes a pending floating-point exception (ES=1) is normally signaled by the processor on the next waiting FP/MMX™ technology instruction. In the following set of circumstances, the exception may be delayed or the FSW register may contain a wrong value:

1. The excepting floating-point instruction is followed by an instruction that accesses memory across a page (4-Kbyte) boundary or its access results in the update of a page table dirty/access bit.
2. The memory accessing instruction is immediately followed by a waiting floating-point or MMX technology instruction.
3. The waiting floating-point or MMX technology instruction retires during a one-cycle window that coincides with a sequence of internal events related to instruction cache line eviction.

Implication: The floating-point exception will not be signaled until the next waiting floating-point/MMX technology instruction. Alternatively it may be signaled with the wrong TOS and condition code values. This erratum has not been observed in any commercial software applications.

Workaround: None identified

Status: For the stepping affected see the *Summary of Changes* at the beginning of this section.

B58. Race Conditions May Exist on Thermal Sensor SMBus Collision Detection/Arbitration Circuitry

Problem: In certain SMBus configurations, when the thermal sensor is used in “hard wired alert” mode along with at least one other device on the bus, the thermal sensor may continue to send its address after losing a collision arbitration in response to an Alert Response Address (ARA) by the SMBus controller.

In order for this erratum to occur, all of the following conditions must be present:

1. The thermal sensor must be configured with alert enabled (default setting).
2. There must be one or more other devices on the SMBus along with the thermal sensor.
3. One or more of these other devices must be also configured with the alert enabled.
4. One or more of these other devices must have a lower address (higher priority) than the thermal sensor.
5. The thermal sensor must generate an SM alert while at least one other device has an SM alert pending to be serviced.

In this situation, the thermal sensor will continue to send its address on the SMBus even if it has a lower priority than the pending alert. When this occurs, the SMBus controller cannot correctly interpret the device address. This may cause the thermal sensor’s alert flag not to clear and may result in SMBus lockup.

Implication: The SMBus controller may see an invalid address and the resulting response of the SMBus controller will vary from implementation to implementation.

Workaround: Remove any one of the five conditions listed above or:

1. In software, use polling mode for the thermal sensor data collection with alert disabled. This software workaround has been validated on both Intel’s test platforms as well as on certain OEM systems.
2. Ensure that the thermal sensor alert may be cleared by a hardware or software mechanism. The implementation of this workaround will be system dependent.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B59. Intermittent Power-on Failure due to Uninitialized Processor Internal Nodes

Problem: If there is no clock source supplied to the processor’s PICCLK pin, the processor may drive an incorrect address for the reset vector at power-on due to uninitialized processor internal nodes. In this scenario when ADS# is asserted, it is possible that the processor drives either the SMI or NMI vector addresses, rather than the reset vector address.

Implication: Systems that provide a clock to the processor’s PICCLK pin are unaffected by this issue. On a system implementation with no clock source supplied to the processor’s PICCLK pin, a small percentage of the systems may intermittently fail to boot, or may fail to resume from a STR or STD state. On the next power-on, the system will likely boot normally.

Workaround: Supply a clock source to the processor’s PICCLK pin.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B60. Selector for the LTR/LLDT Register May Get Corrupted

Problem: The internal selector portion of the respective register (TR, LDTR) may get corrupted if, during a small window of LTR or LLDT system instruction execution, the following sequence of events occur:

1. Speculative write to a segment register that might follow the LTR or LLDT instruction
2. The read segment descriptor of LTR/LLDT operation spans a page (4 Kbytes) boundary; or causes a page fault

Implication: Incorrect selector for LTR, LLDT instruction could be used after a task switch.

Workaround: Software can insert a serializing instruction between the LTR or LLDT instruction and the segment register write.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B61. INIT Does Not Clear Global Entries in the TLB

Problem: INIT may not flush a TLB entry when:

1. The processor is in protected mode with paging enabled and the page global enable flag is set (PGE bit of CR4 register)
2. G bit for the page table entry is set
3. TLB entry is present in TLB when INIT occurs

Implication: Software may encounter unexpected page fault or incorrect address translation due to a TLB entry erroneously left in TLB after INIT.

Workaround: Write to CR3, CR4 or CR0 registers before writing to memory early in BIOS code to clear all the global entries from TLB.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B62. VM bit will be Cleared on a Double Fault Handler

Problem: Following a task switch to a Double Fault Handler that was initiated while the processor was in virtual-8086 (VM86) mode, the VM bit will be incorrectly cleared in EFLAGS.

Implication: When the OS recovers from the double fault handler, the processor will no longer be in VM86 mode.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B63. Memory Aliasing with Inconsistent A and D bits May Cause Processor Deadlock

Problem: In the event that software implements memory aliasing by having two Page Directory Entries (PDEs) point to a common Page Table Entry (PTE) and the accessed and dirty bits for the two PDEs are allowed to become inconsistent the processor may become deadlocked.

Implication: This erratum has not been observed with commercially available software.

Workaround: Software that needs to implement memory aliasing in this way should manage the consistency of the accessed and dirty bits

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B64. Use of Memory Aliasing with Inconsistent Memory Type May Cause System Hang

Problem: Software that implements memory aliasing by having more than one linear addresses mapped to the same physical page with different cache types may cause the system to hang. This would occur if one of the addresses is non-cacheable used in code segment and the other a cacheable address. If the cacheable address finds its way in instruction cache, and non-cacheable address is fetched in IFU, the processor may invalidate the non-cacheable address from the fetch unit. Any micro-architectural event that causes instruction restart will expect this instruction to still be in fetch unit and lack of it will cause system hang.

Implication: This erratum has not been observed with commercially available software.

Workaround: Although it is possible to have a single physical page mapped by two different linear addresses with different memory types, Intel has strongly discouraged this practice as it may lead to undefined results. Software that needs to implement memory aliasing should manage the memory type consistency.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B65. Processor may Report Invalid TSS Fault Instead of Double Fault During Mode C Paging

Problem: When an operating system executes a task switch via a Task State Segment (TSS) the CR3 register is always updated from the new task TSS. In the mode C paging, once the CR3 is changed the processor will attempt to load the PDPTs. If the CR3 from the target task TSS or task switch handler TSS is not valid then the new PDPT will not be loaded. This will lead to the reporting of invalid TSS fault instead of the expected Double fault.

Implication: Operating systems that access an invalid TSS may get invalid TSS fault instead of a Double fault.

Workaround: Software needs to ensure any accessed TSS is valid.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B66. Machine Check Exception may Occur When Interleaving Code Between Different Memory Types

Problem: A small window of opportunity exists where code fetches interleaved between different memory types may cause a machine check exception. A complex set of micro-architectural boundary conditions is required to expose this window.

Implication: Interleaved instruction fetches between different memory types may result in a machine check exception. The system may hang if machine check exceptions are disabled. Intel has not observed the occurrence of this erratum while running commercially available applications or operating systems.

Workaround: Software can avoid this erratum by placing a serializing instruction between code fetches which span different memory types.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B67. Wrong ESP Register Values During a Fault in VM86 Mode

Problem: At the beginning of the IRET instruction execution in VM86 mode, the lower 16 bits of the ESP register are saved as the old stack value. When a fault occurs, these 16 bits are moved into the 32-bit ESP, effectively clearing the upper 16 bits of the ESP.

Implication: This erratum has not been observed to cause any problems with commercially available software.

Workaround: None identified

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B68. APIC ICR Write May Cause Interrupt Not to be Sent When ICR Delivery Bit Pending

Problem: If the APIC ICR (Interrupt Control Register) is written with a new interrupt command while the Delivery Status bit from a previous interrupt command is set to '1' (Send Pending), the interrupt message may not be sent out by the processor.

Implication: This erratum will cause an interrupt message not to be sent, potentially resulting in system hang.

Workaround: Software should always poll the Delivery Status bit in the APIC ICR and ensure that it is '0' (Idle) before writing a new value to the ICR.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B1AP. APIC Access to Cacheable Memory Causes SHUTDOWN

Problem: APIC operations which access memory with any type other than uncacheable (UC) are illegal. If an APIC operation to a memory type other than UC occurs and Machine Check Exceptions (MCEs) are disabled, the processor will enter shutdown after such an access. If MCEs are enabled, an MCE will occur. However, in this circumstance, a second MCE will be signaled. The second MCE signal will cause the Mobile Pentium II processor to enter shutdown.

Implication: Recovery from a PIC access to cacheable memory will not be successful. Software that accesses only UC type memory during APIC operations will not encounter this erratum.

Workaround: Ensure that the memory space to which PIC accesses can be made is marked as type UC (uncacheable) in the memory type range registers (MTRRs) to avoid this erratum.

Status: For the steppings affected see the *Summary of Changes* at the beginning of this section.

B69. Processor Incorrectly Samples NMI Interrupt after RESET# Deassertion When Processor APIC is Hardware-Disabled

Problem: When the processor APIC is hardware-disabled the processor may incorrectly interpret the NMI signal as an NMI interrupt, instead of a frequency strap value, starting six bus clocks after RESET# is deasserted. This will result in a processor hang due to the NMI Handler not being installed at this time.

Implication: The system may fail to boot due to this issue.

Workaround: The processor APIC must be hardware-enabled by pulling PICD[1:0] high with separate pull up resistors and supplying PICCLK to the processor.

Status: For the steppings affected, see the *Summary of Changes* at the beginning of this section.

B70. The Instruction Fetch Unit (IFU) May Fetch Instructions Based Upon Stale CR3 Data After a Write to CR3 Register

Problem - Under a complex set of conditions, there exists a one clock window following a write to the CR3 register where-in it is possible for the iTLB fill buffer to obtain a stale page translation based on the stale CR3 data. This stale translation will persist until the next write to the CR3 register, the next page fault or execution of a certain class of instructions including RDTSC, CPUID, or IRETD with privilege level change.

Implication - The wrong page translation could be used leading to erroneous software behavior.

Workaround - Operating systems that are potentially affected can add a second write to the CR3 register.

Status - For the steppings affected, see the *Summary of Changes* at the beginning of this section.

DOCUMENTATION CHANGES

The Documentation Changes listed in this section apply to the following documents:

- *Mobile Pentium® II Processor at 233 MHz, 266 MHz, and 300 MHz* datasheet
- *Intel® Pentium® II Processor Mobile Module : Mobile Module Connector 1 (MMC-1)* datasheet
- *Intel® Pentium® II Processor Mobile Module : Mobile Module Connector 2 (MMC-2)* datasheet
- *Mobile Pentium® II Processor in Mini-Cartridge Package at 400 MHz, 366 MHz, 333 MHz, 300PE MHz, and 266PE MHz* datasheet
- *Mobile Pentium® II Processor in Micro-PGA and BGA Packages at 400 MHz, 366 MHz, 333 MHz, 300PE MHz, and 266PE MHz* datasheet
- *Intel® Pentium® II Processor Mobile Module With On-Die Cache: Mobile Module Connector 1 (MMC-1)* datasheet
- *Intel® Pentium® II Processor Mobile Module With On-Die Cache: Mobile Module Connector (MMC-2)* datasheet
- *Intel Architecture Software Developer's Manual, Volumes 1, 2, and 3.*
- *P6 Family of Processors Hardware Developer's Manual*

All Documentation Changes will be incorporated into a future version of the appropriate Mobile Intel Pentium II processor documentation.

B1. Misabeled Vertical Dimension Parameters

Figure 5.6 in the *Mobile Pentium® II Processor at 233 MHz, 266 MHz, and 300 MHz* datasheet should be replaced by the following figure:




Figure 5.6. Cross Sectional View with Vertical Dimensions

B2. Correction to Dimension Definitions

In Table 5.3 of the *Mobile Pentium® II Processor at 233 MHz, 266 MHz, and 300 MHz* datasheet, the following correction should be made:

Table 5.3. Mobile Pentium® II Processor Dimension Definitions

Symbol	Parameter	Definition
P	Bottom Cover Height	Distance from the bottom side of the processor substrate to the bottom outer surface of the processor

B3. Correction to Vertical Dimensions

Bottom Cover Height (P) and Note 2 in Table 5.2 of the *Mobile Pentium® II Processor at 233 MHz, 266 MHz, and 300 MHz* datasheet should read as follows:

Table 5.2. Mobile Pentium® II Processor Vertical Dimensions^{1, 2}

Symbol	Parameter	Min	Nom	Max	Units	Figure
H	Processor Height	4.40	4.55	4.70	mm	5.6
M	Plug Height	0.91	1.03	1.15	mm	5.6
P	Bottom Cover Height	—	2.54	2.81	mm	5.6
BLT	Bond-Line Thickness	0.01	0.12	0.23	mm	5.6

NOTES:

- See Table 5.3 and Figure 5.6 for the definitions and illustration of processor height, plug height, bottom cover height and bond-line thickness.
- Processor height, plug height, bottom cover height and bond-line thickness are the only monitored vertical dimensions. The other dimensions in Figure 5.6 are dependent on the receptacle height. Refer to the *Mobile Pentium® II Mechanical and Thermal Design Guide* for information on the effect of receptacle mounting technique on these dimensions.

B4. SMBus Data Setup Time

The SMBus data signal has a min setup time (min Tsu: DAT) of 250ns, as published in the *System Management Bus Specification*, Rev 1.0. This parameter is currently specified to 800 ns for the Mobile Pentium II processor (mini-cartridge) and Pentium II Processor Mobile Module products (see table below).

SMBus Data Min Setup

Parameter	Description	Spec	Change to
min Tsu: DAT	min SMBus Data Setup Time	250 ns	800 ns

B5. Intel® Pentium® II Processor Mobile Module Connector 1 (MMC-1) Northbridge Correction

In the *Intel® Pentium® II Processor Mobile Module : Mobile Module Connector 1 (MMC-1)* datasheet, the northbridge component was incorrectly identified as the 443BX in the MMC-1 versions of the above datasheets. The correct northbridge component for MMC-1 is the 443DX.

B6. VR_ON and VRPWRGD Specification Correction for MMC-1 and MMC-2

In the *Intel® Pentium® II Processor Mobile Module : Mobile Module Connector 1 (MMC-1)* datasheet and the *Intel® Pentium® II Processor Mobile Module : Mobile Module Connector 2 (MMC-2)* datasheet, the VR_ON and VR_PWRGD signal specifications are incorrectly identified as V_3S in all of the above datasheets. The signal specifications should be pulled up to V_3.

B7. STR and POS Measurement Correction for MMC-1 and MMC-2

In the *Intel® Pentium® II Processor Mobile Module : Mobile Module Connector 1 (MMC-1)* datasheet and the *Intel® Pentium® II Processor Mobile Module : Mobile Module Connector 2 (MMC-2)* datasheet, the POS and STR values are incorrectly shown as measured in Amps in all of the above datasheets. The POS and STR values should be measured in Watts. The actual number values are correct.

B8. L2 Cache Correction for MMC-1 and MMC-2

In the *Intel® Pentium® II Processor Mobile Module : Mobile Module Connector 1 (MMC-1)* datasheet and the *Intel® Pentium® II Processor Mobile Module : Mobile Module Connector 2 (MMC-2)* datasheet, the affected datasheets, V_3S is incorrectly identified as the power source for the L2 Cache circuit. V_3S is not the power source for the L2 Cache circuit.

B9. Die Height (A₂) in Surface-Mount BGA1 Package Specification

In the *Mobile Pentium® II Processor in BGA Package at 366 MHz, 333 MHz, 300PE MHz, and 266PE MHz* datasheet, the original documentation of Min at 1.23 mm and Max at 1.38 mm in Table 5.1 is incorrect. The correct documentation is 0.841 mm REF. Refer to the following table for the updated specification.

Table 5.1 Surface-Mount BGA1 Package Specifications

Symbol	Parameter	Min	Max	Unit
A	Overall Height, as delivered	2.29	2.79	mm
A ₁	Ball Height, as delivered	0.76	1.10	mm
A ₂	Die Height	0.841 REF		mm
b	Ball Diameter	0.78 REF		mm

Table 5.1 Surface-Mount BGA1 Package Specifications

Symbol	Parameter	Min	Max	Unit
D	Package Width	30.850	31.150	mm
D ₁	Die Width	10.36 REF		mm
E	Package Length	34.850	35.150	mm
e	Ball Pitch	1.270		mm
E ₁	Die Length	17.36 REF		mm
K	Keepout Outline to Edge of Substrate	5 REF		mm
K ₁	Keepout Outline to Edge of Substrate at Corner	7 REF		mm
N	Ball Count	615		each
S ₁	Outer Ball Center to Short Edge of Substrate	1.625 REF		mm
S ₂	Outer Ball Center to Long Edge of Substrate	0.895 REF		mm
P _{DIE}	Allowable Pressure on the Die for Thermal Solution	—	689	kPa
W	Package Weight	3.71	4.52	grams

B10. MTRR Initialization Clarification

The following sentence should be added to the end of the first paragraph of Section 9.12.5 of the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*: "The MTRRs must be disabled prior to initialization or modification."

B11. STPCLK# Pin Definition

The *Mobile Pentium® II Processor in BGA Package at 366 MHz, 333 MHz, 300PE MHz, and 266PE MHz* datasheet (Order Number 245106) has an incorrect definition of the STPCLK# pin. This document incorrectly states,

The processor continues to snoop bus transactions and *service interrupts* while in Stop Grant state. When STPCLK# is deasserted, the processor restarts its internal clock to all units and resumes execution.

It should state,

The processor continues to snoop bus transactions and *may latch interrupts* while in Stop Grant state. When STPCLK# is deasserted, the processor restarts its internal clock to all units, resumes execution, and *services any pending interrupts*.

B12. RESET# Pin Definition

The *Mobile Pentium® II Processor in BGA Package at 366 MHz, 333 MHz, 300PE MHz, and 266PE MHz* datasheet (Order Number 245106) has an incorrect definition of the RESET# pin. This document incorrectly states,

RESET# must remain active for one microsecond for a 'warm' Reset; for a Power-on Reset, RESET# must stay active for at least one millisecond after VCCCORE and CLK have reached their proper specifications.

It should state,

For a Power-on Reset, RESET# must stay active for at least one millisecond after VCCCORE and CLK have reached their proper specifications.

B13. BGA1 Package Mechanical Specifications

The following table and figures show the updated mechanical specifications for the BGA1 package, listed as Table 5.2, Figure 5.3 and Figure 5.4, respectively, in the *Mobile Pentium® II Processor in BGA Package at 366 MHz, 333 MHz, 300PE MHz, and 266PE MHz* datasheet.

Table 5.2 Surface-Mount BGA1 Package Specifications

Symbol	Parameter	Min	Max	Unit
A	Overall Height, as delivered	2.29	2.79	mm
A ₁	Substrate Height, as delivered	1.50 REF		mm
A ₂	Die Height	0.854 REF ¹		mm
B	Ball Diameter	0.78 REF		mm
D	Package Width	30.85	31.15	mm
D ₁	Die Width	10.36 REF		mm
E	Package Length	34.85	35.15	mm
e	Ball Pitch	1.27		mm
E ₁	Die Length	17.36 REF		mm
N	Ball Count	615		each
S ₁	Outer Ball Center to Short Edge of Substrate	1.625 REF		mm
S ₂	Outer Ball Center to Long Edge of Substrate	0.895 REF		mm
P _{DIE}	Allowable Pressure on the Die for Thermal Solution	—	689	kPa
W	Package Weight	3.71	4.52	grams

NOTE:

1. The number shown here is different from that in the 243887-010 revision of this specification update. The maximum reference number is used.




Figure 5.3. Surface-Mount BGA1 Package—Top and Side View




Figure 5.4. Surface-Mount BGA1 Package—Bottom View

B14. Generic Clock Waveform Figure

Figure 3.3 in *Mobile Pentium® II Processor in Mini-Cartridge Package at 366 MHz, 333 MHz, 300PE MHz, and 266PE MHz* datasheet (Order Number 245108), and the *Mobile Pentium® II Processor in BGA Package at 366 MHz, 333 MHz, 300PE MHz, and 266PE MHz* datasheet (Order Number 245106) has been updated. The following figure is to replace the ones in the above documents.

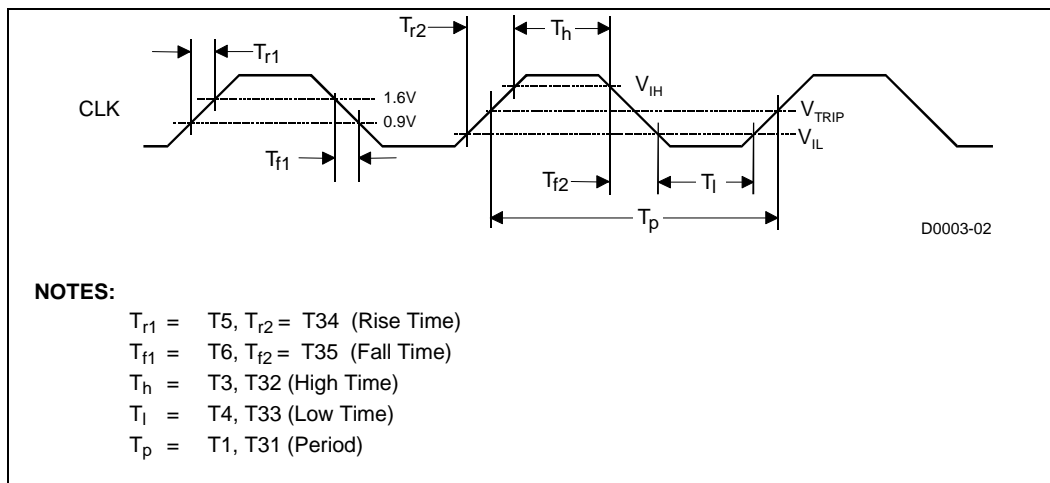


Figure 3.3 Generic Clock Waveform

B15. Invalidating Caches and TLBs

Section 2.6.4 of the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, incorrectly states:

The INVD (invalidate cache with no writeback) instruction invalidates all data and instruction entries in the internal caches and TLBs and sends a signal to the external caches indicating that they should be invalidated also.

It should state:

The INVD (invalidate cache with no writeback) instruction invalidates all data and instruction entries in the internal caches and sends a signal to the external caches indicating that they should be invalidated also.

B16. Handling of Self-Modifying and Cross-Modifying Code

Section 7.1.3 paragraph 1. of the *Intel Architecture Software Developer's Manual Vol 3: System Programming* incorrectly states:

The act of a processor writing data into a currently executing code segment with the intent of executing that data as code is called **self-modifying code**. Intel Architecture processors exhibit model-specific behavior when executing self-modified code, depending upon how far ahead of the current execution pointer the code has been modified. As processor architectures become more complex and start to speculatively execute code ahead of the retirement point (as in the P6 family processors), the rules regarding which code should execute, pre- or post-modification, become blurred. To write self-modifying code and ensure that it is compliant with current and future Intel Architectures one of the following two coding options **should** be chosen.

It should state:

"The act of a processor writing data into a currently executing code segment with the intent of executing that data as code is called **self-modifying code**. Intel Architecture processors exhibit model-specific behavior when executing self-modified code, depending upon how far ahead of the current execution pointer the code has been modified. As processor architectures become more complex and start to speculatively execute code ahead of the retirement point (as in the P6 family processors), the rules regarding which code should execute, pre- or post-modification, become blurred. To write self-modifying code and ensure that it is compliant with current and future Intel Architectures one of the following two coding options **must** be chosen."

Section 7.1.3 paragraph 6. of the *Intel Architecture Software Developer's Manual Vol 3: System Programming* incorrectly states:

The act of one processor writing data into the currently executing code segment of a second processor with the intent of having the second processor execute that data as code is called **cross-modifying code**. As with self-modifying code, Intel Architecture processors exhibit model-specific behavior when executing cross-modifying code, depending upon how far ahead of the executing processors current execution pointer the code has been modified. To write cross-modifying code and insure that it is compliant with current and future Intel Architectures, the following processor synchronization algorithm **should** be implemented.

It should state:

The act of one processor writing data into the currently executing code segment of a second processor with the intent of having the second processor execute that data as code is called **cross-modifying code**. As with self-modifying code, Intel Architecture processors exhibit model-specific behavior when executing cross-modifying code, depending upon how far ahead of the executing processors current execution pointer the code has been modified. To write cross-modifying code and insure that it is compliant with current and future Intel Architectures, the following processor synchronization algorithm **must** be implemented.

B17. Machine Check Architecture Initialization of MCI_STATUS Registers

Section 12.5, the last paragraph of the *Intel Architecture Software Developer's Manual Vol. 3: System Programming Guide* incorrectly states:

The processor can write valid information (such as an ECC error) into the MCI_STATUS registers while it is being powered up. As part of the initialization of the MCE exception handler, software might examine all the MCI_STATUS registers and log the contents of them, then rewrite them all to zeroes. This procedure is not included in the initialization pseudocode in Example 12-1.

It should state:

The processor can write valid information (such as an ECC error) into the MCI_STATUS registers while it is being powered up. As part of the initialization of the MCE exception handler, software might examine all the MCI_STATUS registers and log the contents of them, then rewrite them all to zeroes. Following power cycling, the MCI_STATUS registers are not guaranteed to have valid data until after the registers are initially cleared to all zeroes by software. This procedure is not included in the initialization pseudocode in Example 12-1.

B18. LOCK# Signal Prefix Operands

Page 3-273, the first sentence of the third paragraph of the *Intel Architecture Software Developer's Manual, Volume 2: Instruction Set Reference* states:

The LOCK prefix can be prepended only to the following instructions and to those forms of the instructions that use a memory operand: ADD, ADC, AND, ...

It should state:

The LOCK prefix can be prepended only to the following instructions and to those forms of the instructions that use a memory operand as the destination operand only: ADD, ADC, AND, ...

If the LOCK prefix is used with the memory operand as the source operand then an Invalid Opcode (Undefined Opcode), #UD, can occur.

B19. SMRAM State Save Map Contains Documentation Errors

In the *Intel Architecture Software Developer's Manual Volume 3: System Programming Guide*, revision 2 Chapter 12, "System Management Mode (SMM)," Table 12-1 incorrectly documents the SMBASE+Offset for LDT Base on the P6 family of processors.

The storage locations for these parameters are model specific (i.e., they may differ between the Pentium® processor, the Pentium® Pro processor, and other P6 family proliferations). **These entries in the tables above will be changed to Reserved. Hardware and software may not rely on the contents of these Reserved regions.**

System Management Interrupt (SMI) During Startup IPI

B20. Clarification

The note on section 12.2 of Intel Architecture Software Developer's Manual Vol. 3: System Management Interrupt (SMI) states:

In the P6 families of processors, when a processor that is designated as the application processor during an MP initialization protocol is waiting for a startup IPI (SIPI), it is in a mode where SMIs are masked.

It should state:

In the P6 family of processors, when a processor that is designated as the application processor during an MP initialization protocol is waiting for a startup IPI (SIPI), it is in a mode where SMIs are masked. However if SMI is received while an application processor is in the wait for SIPI mode it will be pended. The processor will respond on receipt of a SIPI by immediately servicing the pended SMI and will go into SMM before executing from the SIPI vector.

B21. Memory Aliasing with Different Memory Types

The 4th paragraph of section 9.13.5 of Intel Architecture Software Developer's Manual Vol. 3: Programming the PAT states:

The PAT allows any memory type to be specified in the page table, and therefore it is possible to have a single physical page mapped to two different linear addresses with different memory types. This practice is strongly discouraged by Intel and should be avoided as it may lead to undefined results.

It should change to:

The PAT allows any memory type to be specified in the page table, and therefore it is possible to have a single physical page mapped to two different linear addresses with different memory types. Intel does not support this practice as it may lead to undefined operations including processor hang.

B22. Runbist will Not Function When stpclk# Driven Low

Paragraph 5 of Section 6.3 in the *P6 Family of Processors Hardware Developer's Manual* (Order Number 244001) currently states:

Note that RUNBIST will not function when the processor core clock has been stopped. All other 1149.1-defined instructions operate independently of the processor core clock.

It should state:

Note that RUNBIST will not function when the processor STPCLK# input has been driven low. All other 1149.1-defined instructions will correctly operate regardless of the STPCLK# signal state.

B23. Memory Aliasing with Inconsistent A and D Bits may Cause Processor Deadlock

Add the following note to Chapter 3 and Chapter 9.13.5 of the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide, Revision 2*:

Processor allows memory aliasing by having two Page Directory Entries (PDEs) point to a common Page Table Entry (PTE). Software that needs to implement memory aliasing in this way should manage the consistency of the Accessed and Dirty bits. Allowing the Accessed and Dirty bits for the two PDEs to become inconsistent may lead to a processor deadlock.

B24. An Interrupt Could Occur While TSS is Marked Busy

Paragraph 5 of section 6.1.3 in the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, currently states:

For all Intel Architecture processors, tasks are not recursive. A task cannot call or jump to itself.

It should state:

For all Intel Architecture processors, tasks are not recursive. A task cannot call or jump to itself. Because Intel Architecture tasks are not re-entrant, a task used as an interrupt handler must have interrupts disabled between the time it completes the interrupt and the time it exits with IRET. Otherwise, another interrupt could occur while the interrupt task's TSS is still marked busy, causing a #GP fault.

NMI Unmasked Early When Processor is Running in V86

B25. Mode

Section 5.5.1 in the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, currently states:

While an NMI interrupt handler is executing, the processor disables additional calls to the NMI handler until the next IRET instruction is executed. This blocking of subsequent NMIs prevents stacking up calls to NMI handler. It is recommended that the NMI interrupt handler be accessed through an interrupt gate to disable maskable hardware interrupts (refer to section 5.6.1, "Masking Maskable Hardware Interrupts").

It should state:

While an NMI interrupt handler is executing, the processor disables additional calls to the NMI handler until the next IRET instruction is executed. This blocking of subsequent NMIs prevents stacking up calls to NMI handler. It is recommended that the NMI interrupt handler be accessed through an interrupt gate to disable maskable hardware interrupts (refer to section 5.6.1, "Masking Maskable Hardware Interrupts"). If the NMI handler is a virtual-8086 task with IOPL less than 3, the IRET from this handler triggers a general-protection exception (#GP) (section 16.2.7). In this case, NMI is unmasked before the #GP handler is invoked.

P6 Family Processors Read Two Bytes for POP SEG B26. Instruction

Paragraph 1 and 2 of section 18.24.1 in the Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide, currently states:

When pushing a segment selector onto the stack, the Intel486(TM) processor writes 2 bytes onto 4-byte stacks and decrements ESP by 4. The P6 family and Pentium® processors write 4 bytes, with the upper 2 bytes being zeros.

When popping a segment selector from the stack, the Intel486(TM) processor reads only 2 bytes. The P6 family and Pentium® processors read 4 bytes and discard the upper 2 bytes. This operation may have an effect if the ESP is close to the stack-segment limit. On the P6 family and Pentium® processors, stack location at ESP plus 4 may be above the stack limit, in which case a stack fault exception (#SS) will be generated. On the Intel486(TM) processor, stack location at ESP plus 2 may be less than the stack limit and no exception is generated.

It should state:

When pushing a segment selector, Intel486(TM) and P6 family processors decrement ESP by the operand size and then write two bytes. If the operand size is 32-bits, the upper two bytes of the write are unmodified. The Intel Pentium® processor decrements ESP by the operand size and determines the size of the write by the operand size. If the operand size is 32-bits, the upper two bytes of the write are zero.

When popping a segment selector, Intel486(TM) and P6 family processors read two bytes and increment ESP by the operand size of the instruction. The Intel Pentium® processor determines the size of the read by the operand size and increments ESP by the operand size.

It is possible to align a 32-bit selector push or pop such that the operation will generate an exception on the Intel Pentium® processor and not on the Intel486(TM) and P6 family processors. This could occur if the third and/or fourth byte of the operation lies beyond the limit of the segment or if the third and/or fourth byte of the operation is located on a not-present or inaccessible page.

B27. APIC Register Offsets are Aligned on 128-bit Boundaries

Paragraph 3 of section 7.5.7 in the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, currently states:

Within the 4KB APIC register area, the register address allocation scheme is shown in Table 7-1. Register offsets are aligned on 128-bit boundaries. All registers must be accessed using 32-bit loads and stores. Wider registers (64-bit or 256-bit) are defined and accessed as independent multiple 32-bit registers. If a LOCK prefix is used with a MOV instruction that accesses the APIC address space, the prefix is ignored; that is a locking operation does not take place.

It should say:

Within the 4KB APIC register area, the register address allocation scheme is shown in Table 7-1. Register offsets are aligned on 128-bit boundaries. All registers must be accessed using loads and stores that are aligned on 128-bit boundaries and manipulate the registers as 32-bit quantities. Wider registers (64-bit or 256-bit) are defined and accessed as independent multiple 32-bit registers. If a LOCK prefix is used with a MOV instruction that accesses the APIC address space, the prefix is ignored; that is a locking operation does not take place.

B28. Single Stepping of Instructions Breaks out of HALT State

Paragraph 1 of section 2.6.5 in the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, currently states:

The HLT (halt processor) instruction stops the processor until an enabled interrupt (such as NMI or SMI, which are normally enabled), the BINIT# signal, the INIT# signal, or the RESET# signal is received. The processor generates a special bus cycle to indicate that the halt mode has been entered. Hardware may respond to this signal in a number of ways. An indicator light on the front panel may be turned on. An NMI interrupt for recording diagnostic information may be generated. Reset initialization may be invoked. (Note that the BINIT# pin was introduced with the Pentium® Pro processor)

It should say:

The HLT (halt processor) instruction stops the processor until an enabled interrupt (such as NMI, or SMI which are normally enabled), a debug exception, the BINIT# signal, the INIT# signal, or the RESET# signal is received. The processor generates a special bus cycle to indicate that the halt mode has been entered. Hardware may respond to this signal in a number of ways. An indicator light on the front panel may be turned on. An NMI interrupt for recording diagnostic information may be generated. Reset initialization may be invoked. (Note that the BINIT# pin was introduced with the Pentium® Pro processor)

B29. Additional Signal Resumes Execution While in a HALT State

Paragraph 1 of Description section of page 3-291 in the *Intel Architecture Software Developer's Manual, Volume 2: Instruction Set Reference Manual*, currently states:

This instruction stops instruction execution and places the processor in a HALT state. An enabled interrupt, NMI, or a reset will resume execution. If an interrupt (including NMI) is used to resume execution after a HLT instruction, the saved instruction pointer (CS:EIP) points to the instruction following the HLT instruction.

It should say:

This instruction stops instruction execution and places the processor in a HALT state. An enabled interrupt (including NMI and SMI), a debug exception, the BINIT# signal, the INIT# signal, or the RESET# signal will resume execution. If an interrupt (including NMI) is used to resume execution after a HLT instruction, the saved instruction pointer (CS:EIP) points to the instruction following the HLT instruction.

B30. PDPTR Loads are Always Uncacheable

Paragraph 1 of section 3.8.4 in the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, currently states:

Figure 3-20 shows the format for the page-directory-pointer-table, page-directory, and page-table entries when 4-KByte pages and 36-bit extended physical addresses are being used. Figure 3-21 shows the format for the page-directory-pointer-table and page-directory entries when 2-MByte pages and 36-bit extended physical addresses are being used. The functions of the flags in these entries are the same as described in Section 3.6.4., "Page-Directory and Page-Table Entries". The major differences in these entries are as follows:

- A page-directory-pointer-table entry is added.
- The size of the entries are increased from 32 bits to 64 bits.
- The maximum number of entries in a page directory or page table is 512.
- The base physical address field in each entry is extended to 24 bits.

It should state

Figure 3-20 shows the format for the page-directory-pointer-table, page-directory, and page-table entries when 4-KByte pages and 36-bit extended physical addresses are being used. Figure 3-21 shows the format for the page-directory-pointer-table and page-directory entries when 2-MByte pages and 36-bit extended physical addresses are being used. The functions of the flags in these entries are the same as described in Section 3.6.4., "Page-Directory and Page-Table Entries". The major differences in these entries are as follows:

- A page-directory-pointer-table entry is added.
- The size of the entries is increased from 32 bits to 64 bits.
- The maximum number of entries in a page directory or page table is 512.
- The base physical address field in each entry is extended to 24 bits.

Note:

Initial processors that implement the PAE address translation mechanism use uncached accesses when loading page-directory-pointer table entries (PDPTRs). This behavior is model-specific and not architectural. Future processors may cache PDPTRs.

SMI During HALT Causes PMC Miscounts of Retired B31. Instructions

Table A-6 of Appendix A in the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, currently states:

Table A-6. Events That Can Be Counted with the P6 Family Performance-Monitoring Counters (Contd.)

Unit	Event Num.	Mnemonic Event Name	Unit Mask	Description	Comments
Instruction Decoding and Retirement	C0H	INST_RETIRED	OOH	Number of instructions retired.	A hardware interrupt received during/after the last iteration of the REP STOS flow causes the counter to undercount by 1 instruction.

It should state:

Table A-6. Events That Can Be Counted with the P6 Family Performance-Monitoring Counters (Contd.)

Unit	Event Num.	Mnemonic Event Name	Unit Mask	Description	Comments
Instruction Decoding and Retirement	C0H	INST_RETIRED	OOH	Number of instructions retired.	<p>A hardware interrupt received during/after the last iteration of the REP STOS flow causes the counter to undercount by 1 instruction.</p> <p>An SMI received while executing a HLT instruction will cause the performance counter to not count the RSM instruction and therefore undercount by 1.</p>

B32. INIT Message is Not Sent Out Through APIC

In section 7.6.13 in the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, currently states:

101 (INIT Level De-assert)

(The trigger mode must also be set to 1 and level mode to 0.) Sends a synchronization message to all APIC agents to set their arbitration IDs to the values of their APIC IDs. Note that the INIT interrupt is sent to all agents, regardless of the destination field value. However, at least one valid destination processor should be specified. For future compatibility, the software is requested to use a broadcast-to-all ("all-incl-self" shorthand, as described below).

It should state:

101 (INIT Level De-assert)

(The trigger mode must also be set to 1 and level mode to 0.) Sends a synchronization message to all APIC agents to set their arbitration IDs to the values of their APIC IDs. Note that the INIT interrupt is sent to all agents, regardless of the destination field value; however, at least one valid destination processor should be specified. For future compatibility, the software should use the "all excluding self" shorthand instead of designating a destination processor.

B33. Mobile Pentium II Processor CUID Section Update

The Mobile Intel® Pentium® II Processor in BGA2 and Micro-PGA2 Packages datasheet (order number 243669-002) has the following Documentation Change. The datasheet was incorrectly documented as stating the EAX register contained the values after a power-on RESET. The correct statement is after a power-on RESET, the EDX register contains the processor version information (type, family, model, stepping).

Section 2.4 Update.

When the CUID version information is loaded with EAX=01H, the EAX register contain the values shown in Table 2.4. After a power-on RESET, the EDX register contains the processor version information (type, family, model, stepping). See Intel Processor Identification and the CUID Instruction Application Note AP-485 for further information.

Table 2.4. Mobile Pentium II Processor CUID

EAX[31:0]				
Reserved [31:14]	Type [13:12]	Family [11:8]	Model [7:4]	Stepping [3:0]
X	0	6	5	X

SPECIFICATION CLARIFICATIONS

The Specification Clarifications listed in this section apply to the following documents:

- *Mobile Pentium® II Processor at 233 MHz, 266 MHz, and 300 MHz* datasheet
- *Intel® Pentium® II Processor Mobile Module : Mobile Module Connector 1 (MMC-1)* datasheet
- *Intel® Pentium® II Processor Mobile Module : Mobile Module Connector 2 (MMC-2)* datasheet
- *Mobile Pentium® II Processor in Mini-Cartridge Package at 400 MHz, 366 MHz, 333 MHz, 300PE MHz, and 266PE MHz* datasheet
- *Mobile Pentium® II Processor in Micro-PGA and BGA Packages at 400 MHz, 366 MHz, 333 MHz, 300PE MHz, and 266PE MHz* datasheet
- *Intel® Pentium® II Processor Mobile Module With On-Die Cache: Mobile Module Connector 1 (MMC-1)* datasheet
- *Intel® Pentium® II Processor Mobile Module With On-Die Cache: Mobile Module Connector (MMC-2)* datasheet
- *P6 Family of Processors Hardware Developer's Manual*

All Specification Clarifications will be incorporated into a future version of the appropriate Mobile Pentium II processor documentation.

B1. PWRGOOD Inactive Pulse Width

Footnote 10 of Table 3.11 in the *Mobile Pentium® II Processor at 233 MHz, 266 MHz, and 300 MHz* datasheet, should read as follows:

10. When driven inactive or after V_{CC} , V_{CCP} , V_{CC3} , and BCLK become stable. PWRGOOD must remain below $V_{IL,max}$ from Table 3.7 until all the voltage planes meet the voltage tolerance specifications in Table 3.4 and BCLK has met the BCLK AC specifications in Table 3.8 for at least 10 clock cycles. PWRGOOD must rise glitch-free and monotonically to 2.5 V.

B2. Thermal Sensor Setpoint

The thermal sensor in the Mobile Pentium II processor and mobile module implements the SMBALERT# signal described in the SMBus specification. SMBALERT# is always asserted when the temperature of the processor core thermal diode or the thermal sensor internal temperature exceeds either the upper or lower temperature thresholds. SMBALERT# may also be asserted if the measured temperature equals either the upper or lower threshold.

B3. Thermal Sensor Configuration Register—RUN/STOP Bit

(Refer to *Mobile Pentium® II Processor and Pentium® II Processor Mobile Module Thermal Sensor Interface Specifications*, Rev.1.0. AP-825)

The configuration register of the thermal sensor in Mobile Pentium II processor based systems controls the operating mode (Auto-convert vs. Standby) of the device. Since the processor temperature varies dynamically during normal operation, auto-convert mode should be used exclusively to monitor processor temperature. Table 1 shows the format of the configuration register. If the RUN/STOP bit is low, then the thermal sensor enters auto-conversion mode. If the RUN/STOP bit is set high, then the thermal sensor immediately stops converting and enters Standby mode. The thermal sensor will still perform temperature conversions in Standby mode when it receives a one-shot command. However, the result of a one-shot command during auto-convert mode is not guaranteed. Intel does not recommend using the one-shot command to monitor temperature when the processor is active—only auto-convert mode should be used.

Table 1. Thermal Sensor Configuration Register

Bit	Name	Reset State	Function
7 (MSB)	MASK	0	Masks SMBALERT# when high.
6	RUN/STOP	0	Standby mode control bit. If low, the device enters auto-convert mode. If high, the device immediately stops converting, and enters standby mode where the one-shot command can be performed.
5 – 0	RFU	0	Reserved for future use.

NOTE:

All RFU bits should be written as “0” and read as “don’t care” for programming purposes.

B4. Thermal Attach Point Gageable Thread Clarification

In the *Intel® Pentium® II Processor Mobile Module : Mobile Module Connector 1 (MMC-1)* and the *Intel® Pentium® II Processor Mobile Module : Mobile Module Connector 2 (MMC-2)* datasheets (Order Numbers 243667 and 243668, respectively), the thermal attach point (TAP) gageable thread specification was not clearly identified in Figure 14, of Section 5.2. The following is the TAP clarification:

OEM Thermal Attach Points (TAPs) M2.0 x 0.4, 2.25 min gageable thread dimensioned per ANSI specification.

See Figure 1 below for more information.




Figure 1. TAP Gageable Thread Clarification

B5. Shipping Container Maximum Temperature Rating for BGA1

In the *Mobile Pentium® II Processor in BGA Package at 366 MHz, 333 MHz, 300PE MHz, and 266PE MHz* datasheet, note 1 is added to Table 3.5 to include the maximum shipping container temperature rating of 65° C. The following is the updated Table 3.5.

Table 3.5. Mobile Celeron™ Processor Absolute Maximum Ratings

Parameter	Min	Max	Unit	Notes
Storage Temperature	–40	85	°C	1
Supply Voltage with respect to V_{SS}	–0.5	3.0	V	
CMOS Reference Voltage with respect to V_{SS}	–0.3	3.0	V	
GTL+ Buffer DC Input Voltage with respect to V_{SS}	–0.3	$V_{CC} + 0.7$	V	2
2.5 V Buffer DC Input Voltage with respect to V_{SS}	–0.3	3.3	V	3

NOTES:

1. The shipping container is only rated for 65° C.
2. Parameter applies to the Low Power GTL+ signal groups only.
3. Parameter applies to CMOS, Open-Drain, APIC and TAP bus signal groups only.

B6. ESD for MMC-1 and MMC-2

In the *Intel® Pentium® II Processor Mobile Module : Mobile Module Connector 1 (MMC-1)* datasheet and the *Intel® Pentium® II Processor Mobile Module : Mobile Module Connector 2 (MMC-2)* datasheet, ESD is a nonpowered test of the module for noncatastrophic failure only. The module is tested at 2 KV and then is inserted in a system for a functional test.

B7. Thermal Diode Footnote

In the *Mobile Pentium® II Processor in Micro-PGA and BGA Packages at 400 MHz, 366 MHz, 333 MHz, 300PE MHz, and 266PE MHz* datasheet, footnote 2 for Table 6.3 states,

“2. For the 400 MHz processor. Characterized at 100° C.”

It should state as

“2. For the 400 MHz processor. Characterized at 100° C with forward bias current from 5 to 300 μ A.”

B8. Bulk Capacitance Requirements for the System Electronics

This is a clarification to the *Intel® Pentium® II Processor Mobile Module : Mobile Module Connector 1 (MMC-1) datasheet (Order Number 243667)*, *Intel® Pentium® II Processor Mobile Module : Mobile Module Connector 2 (MMC-2) datasheet (Order Number 243668)*, *Pentium® II Processor Mobile Module with on-die Cache, MMC-1 datasheet (Order Number 245109)*, *Pentium® II Processor Mobile Module with on-die Cache, MMC-2 datasheet (Order Number 245110)*:

Placement of sufficient bulk capacitance on the system electronics board is critical to the operation of the Intel Pentium II processor mobile module and to ensure the system design can accommodate future high frequency modules. Intel has provided the maximum possible bulk capacitance on the module. However, in order to achieve proper filtering and in-rush current protection, it is imperative that additional filtering be provided on the system electronics board. Table 1 details the bulk capacitance requirements for the system electronics. Also, take special note of the voltage rating requirement for the capacitors on each respective voltage rail.

Note that the rating requirement and note 5 has been changed from previous versions of this table in the *Pentium® II Processor Mobile Module datasheet*. Future revisions of this specification will include these new changes.

Table 2. Bulk Capacitance

Power Plane	Capacitance Requirements	Max Total ESR	Ripple Current	Rating
V_DC	100 uf ⁵ , 0.1 uf, 0.01 uf ¹	20 mΩ	4.5A-7.5A ³ 3A ~ 5A ⁴	20% tolerance at 35V
V_5	100 uf ⁵ , 0.1 uf, 0.01 uf ¹	100 mΩ	1A	20% tolerance at 10V
V_3	470 uf ⁵ , 0.1 uf, 0.01 uf ¹	100 mΩ	1A	20% tolerance at 6V
V_3S	100 uf ⁵ , 0.1 uf, 0.01 uf ¹	100 mΩ	N/A	20% tolerance at 6V
VCC_AGP	22 uf ⁵ , 0.1 uf, 0.01 uf ¹	100 mΩ	1A	20% tolerance at 6V
V_CPUPU	2.2 uf, 8200 pf ¹	N/A	N/A	20% tolerance at 6V
V_CLK	10 uf ⁵ , 8200 pf ²	N/A	N/A	20% tolerance at 6V

NOTES:

1. Placement of above capacitance requirements should be located near the connector.
2. V_CLK filtering should be located next to the system clock synthesizer.
3. Ripple current specification depends on V_DC input for the Geyserville module.
4. Ripple current specification depends on V_DC input for the single frequency module.
5. If Tantalum Capacitors are used, a 50% voltage de-rating practice must be observed (for example, a 5.0-V rail requires a 10.0-V rated capacitor).

SPECIFICATION CHANGES

The Specification Changes listed in this section apply to the following documents:

- *Mobile Pentium® II Processor at 233 MHz, 266 MHz, and 300 MHz* datasheet
- *Intel® Pentium® II Processor Mobile Module : Mobile Module Connector 1 (MMC-1)* datasheet
- *Intel® Pentium® II Processor Mobile Module : Mobile Module Connector 2 (MMC-2)* datasheet
- *Mobile Pentium® II Processor in Mini-Cartridge Package at 400 MHz, 366 MHz, 333 MHz, 300PE MHz, and 266PE MHz* datasheet
- *Mobile Pentium® II Processor in Micro-PGA and BGA Packages at 400 MHz, 366 MHz, 333 MHz, 300PE MHz, and 266PE MHz* datasheet
- *Intel® Pentium® II Processor Mobile Module With On-Die Cache: Mobile Module Connector 1 (MMC-1)* datasheet
- *Intel® Pentium® II Processor Mobile Module With On-Die Cache: Mobile Module Connector (MMC-2)* datasheet
- *P6 Family of Processors Hardware Developer's Manual*

All Specification Changes will be incorporated into a future version of the appropriate Mobile Pentium II processor documentation.

B1. Maximum Die and Cover Pressure Specifications

The die and top cover pressure specification in Table 5.1 in the *Mobile Pentium® II Processor at 233 MHz, 266 MHz, and 300 MHz* datasheet should be changed as follows:

Table 5.1. Mobile Pentium® II Processor Mechanical Specifications

Symbol	Parameter	Max	Unit	Notes
P _{DIE}	Allowable Pressure On Die For Thermal Attach Plate (TAP)	411	kPa	
P _{TOP}	Allowable Pressure on Recessed Part of the Top Cover	411	kPa	

B2. New Footnote for PWRGOOD Inactive Pulse Width

In Table 3.11 of the *Mobile Pentium® II Processor at 233 MHz, 266 MHz, and 300 MHz* datasheet, the following addition should be made:

Table 3.11. PC Compatibility Signal Group AC Specifications^{1, 2}

Symbol	Parameter	Min	Max	Unit	Figure	Notes
T15	PWRGOOD Inactive Pulse Width	10		BCLKs	3.6	10, 11

Also, the following footnote should be added:

11. If the BCLK signal meets its AC specification within 150 ns of turning on, then the PWRGOOD Inactive Pulse Width specification (T15) is waived and BCLK may start after PWRGOOD is asserted. PWRGOOD must still remain below $V_{IL,max}$ until all the voltage planes meet the voltage tolerance specifications.

B3. APIC Reference Correction

Advanced Programmable Interrupt Controller (APIC) is not a supported mobile feature. References to this feature were inadvertently included in the current datasheet. Please disregard any and all references.

B4. Low Power GTL+ Output Leakage Current Change

The Low Power GTL+ Output Leakage Current specification is changing from $\pm 15 \mu A$ to $\pm 30 \mu A$. Please refer to the following changes applicable to Table 3.7 of the *Intel® Mobile Pentium® II Processor* datasheet:

Table 3.7. Low Power GTL+ Signal Group DC Specifications

Symbol	Parameter	Min	Max	Unit	Notes
I_{LO}	Output Leakage Current		± 30	μA	1

NOTES:

1. $(0 \leq V_{OUT} \leq 2.5 V + 5\%)$.

B5. $I_{CC,SG}$, $I_{CC,QS}$, and $I_{CC,DSL P}$ Maximum Specifications for BGA1 and Mini-Cartridge

In the Mobile Pentium® II Processor in Mini-Cartridge Package at 366 MHz, 333 MHz, 300PE MHz, and 266PE MHz datasheet and the Mobile Pentium® II Processor in BGA Package at 366 MHz, 333 MHz, 300PE MHz, and 266PE MHz datasheet, maximum $I_{CC,SG}$, $I_{CC,QS}$, and $I_{CC,DSL P}$ specifications have been increased by 250 mA, to 1190 mA, 880 mA and 650 mA, respectively. The following is the updated Power Specification table of the datasheets.

Table 3.6. Mobile Celeron™ Processor Power Specifications¹

$T_{CASE} = 0$ to $T_{CASE,max}$; $V_{CC} = 1.6\text{ V} \pm 135\text{mV}$; $V_{CCP} = 1.8\text{ V} \pm 90\text{mV}$						
Symbol	Parameter	Min	Typ	Max	Unit	Notes
V_{CC}	V_{CC} of core logic for regular voltage processors	1.465	1.6	1.735	V	$\pm 135\text{ mV}$
$V_{CC,LP}$	V_{CC} when $I_{CC} < 300\text{ mA}$	1.465	1.6	1.805	V	$+205/-135\text{ mV}^2$
V_{CCP}	V_{CC} for CMOS voltage references	1.71	1.8	1.89	V	$1.8\text{ V} \pm 90\text{ mV}$
I_{CC}	I_{CC} for V_{CC} at core frequency @ 336 MHz @ 300 MHz @ 300PE MHz @ 266PE MH			8.80 7.95 7.49 6.63	A A A A	5
I_{CCP}	Current for V_{CCP}			75	mA	3, 4, 5
$I_{CC,SG}$	Processor Stop Grant and Auto Halt current			1190	mA	5
$I_{CC,QS}$	Processor Quick Start and Sleep current			880	mA	5
$I_{CC,DSL P}$	Processor Deep Sleep leakage current			650	mA	5
dI_{CC}/dt	V_{CC} power supply current slew rate			20	A/ μs	6, 7

NOTES:

- Unless otherwise noted, all specifications in this table apply to all processor frequencies.
- A higher $V_{CC,MAX}$ is allowed when the processor is in a low power state to enable high efficiency, low current modes in the power regulator.
- I_{CCP} is the current supply for the CMOS voltage references.
- Not 100% tested. Specified by design/characterization.
- $I_{CCx,max}$ specifications are specified at $V_{CC,max}$, $V_{CCP,max}$ and 100° C and under maximum signal loading conditions. $I_{CCx,max}$ specifications are not specified at $V_{CC,LP,max}$; if that voltage specification is used then slightly higher currents can be expected.
- Based on simulations and averaged over the duration of any change in current. Use to compute the maximum inductance and reaction time of the voltage regulator. This parameter is not tested.
- Maximum values specified by design/characterization at nominal V_{CC} and V_{CCP} .

B6. $I_{CC,SG}$, $I_{CC,QS}$, and $I_{CC,DSL P}$ Maximum Specifications for BGA1 and Mini-Cartridge

In the *Mobile Pentium® II Processor in BGA Package at 366 MHz, 333 MHz, 300PE MHz, and 266PE MHz* datasheet, Ball AC21 was defined as NC (no-connect). This ball has been re-defined as V_{TOL} , which can be used to identify Mobile Pentium II processors in 0.25 μm process from the later generation of Mobile Pentium II processors. In the Mobile Pentium II processors (0.25 μm), this signal is an open-drain signal at high impedance. Unlike other open drain signals, V_{TOL} is 1.6 V tolerant, with a maximum specification of 1.735 V. The following table is a revised Table 3.10 of the datasheet.

Table 3.10. Clock, APIC, TAP, CMOS and Open-Drain Signal Group DC Specifications

$T_{CASE} = 0$ to $T_{CASE,max}$; $V_{CC} = 1.6 V \pm 135 mV$, or $1.5 V \pm 135 mV$; $V_{CCP} = 1.8 V \pm 90 mV$					
Symbol	Parameter	Min	Max	Unit	Notes
V_{IL}	Input Low Voltage	-0.3	0.7	V	
$V_{IL,BCLK}$	Input Low Voltage, BCLK	-0.3	0.7	V	
V_{IH}	Input High Voltage	1.7	2.625	V	
$V_{IH,BCLK}$	Input High Voltage, BCLK	1.8	2.625	V	
V_{OL}	Output Low Voltage		0.4	V	1
V_{OH}	Output High Voltage	N/A	2.625	V	All outputs are open-drain
V_{TOL}	V_{TOL} Voltage	0	1.735	V	
I_{OL}	Output Low Current		14	mA	
I_{LI}	Input Leakage Current		± 100	μA	2
I_{LO}	Output Leakage Current		± 30	μA	2

NOTES:

1. Parameter measured at 14 mA.
2. ($0 \leq V_{IN} \leq 2.625 V$).

B7. Temperature Note for Thermal Power Specification

In the *Mobile Pentium® II Processor in BGA Package at 366 MHz, 333 MHz, 300PE MHz, and 266PE MHz* datasheet, Table 6.1 TDP notes should state “at 100° C^{2, 3}” instead of “at 50° C^{2, 3}”. Table Note 3 should state “3. Not 100% tested or guaranteed. The power specifications are composed of the current of the processor on the various voltage planes. These currents are measured and specified at high temperature in Section 3.5. The 50° C power specifications are determined by characterization of the processor currents at higher temperatures.”